



CODE Presents:

Advances and Challenges in Generative AI for Image Generation

**October 4th, 2023
Online**

Otto Dobretsberger, PhD

Sr. Software Dev at CODE Consulting

CODE
CONSULTING

CODE
STAFFING

CODE
MAGAZINE

CODE
TRAINING



"Helping People Build
Better Software"



About Otto Dobretsberger

CODE
PRESENTS

- Sr. Software Dev @ CODE (15+ years)
- Degrees in Biomedical Science and Computer Science
- Professor @ University of Houston (10+ years)
- Author, Speaker and Presenter

- 2x World Record holder in Bench Press

- [Email: ottod@codemag.com](mailto:ottod@codemag.com)



About CODE Consulting



- Custom Software Development
- Artificial Intelligence
- Web, Cloud, Mobile, Desktop, Serverless, Databases,...
- Copilot Development, AI, GPT & Azure OpenAI
- Application Security Testing
- User Interface and Interaction Design
- Project Rescue, App Modernization (VB, VFP, Access, etc.)
- Support & Maintenance for existing applications



AI Consulting Services



CODE CONSULTING Home Magazine Consulting Staffing Training Framework VFP Conversion [Login](#)

Executive Briefing (AI)

Are you wondering how AI can help your business? Do you worry about privacy or regulatory issues stopping you from using AI to its fullest? We have the answers! Our Executive Briefings provide guidance and concrete advice that help decision makers move forward in this rapidly changing Age of Artificial Intelligence and Copilots!

We will send an expert to your office to meet with you. You will receive:

1. An overview presentation of the current state of Artificial Intelligence.
2. How to use AI in your business while ensuring privacy of your and your clients' information.
3. A sample application built on your own HR documents - allowing your employees to query those documents in English and cutting down the number of questions that you and your HR group have to answer.
4. A roadmap for future use of AI catered to what you do.

A Briefing Combined with Solid Deliverables!

We will send an expert to your office to meet with you. You will receive:

1. An overview presentation of the current state of Artificial Intelligence.
2. How to use AI in your business while ensuring the privacy of your data.
3. A sample application built on your own HR documents - allowing your employees to query those documents in English and cutting down the number of questions that you and your HR group have to answer.
4. A roadmap for future use of AI catered to what you do.

Executive Briefings are set up specifically for you and usually take several hours and up to a day. It is up to you whether you want to include only C-level Executives or your entire company. The cost is **\$5,000** (which is a discounted special rate, compared to the normal price of \$7,500) plus expenses if we have to travel to your location, or stay over night. This includes the concrete deliverable of a first implementation of AI for your business (see #3 above, although we may also agree that something slightly different is more appropriate for your organization).

[Contact us today to set up your individual Executive Briefing!](#)

Setting Directions for the Future of Your Organization

We are experiencing a time of extreme change in technology. While one might argue that this has always been true, it is clear that current developments have a drastic impact on all organizations at a scale that is orders of magnitude beyond what we have experienced in years past.

This represents a time of extreme opportunity, as well as a time of considerable risks!

Artificial Intelligence (AI) continues to revolutionize industries across the globe, offering immense potential for organizations to enhance operational efficiency, drive innovation, and achieve competitive advantage. Conversely, ignoring these technological trends means that even very recent technology is quickly turning into "legacy" and your competitors now have the opportunity to leapfrog you. Especially the advent of "The Age of Copilots" (powered largely by the onset of Large Language Models - LLMs) promises a productivity improvement of upwards of 50% in some areas year after year, and that isn't the only trend to pay attention to.

What does all of this mean specifically for your organization?

- Check out our new **Executive Briefing Offer!**
- We can help with your AI needs
- What does AI mean for you?
- "Skunk Works" Projects



- www.codemag.com/AI
- www.codemag.com/ExecutiveBriefing

CODE Staffing



- Disrupting the world of staffing!
- Giving our customers the ability to have staff on par with Silicon Valley companies...
- ...and our employees a work environment in a bleeding-edge tech company with the **industry leading benefits!**
- www.codestaffing.com



We Need You!

- We are currently “hiring on all cylinders” across the entire CODE Group Organization!
- Check out www.codestaffing.com/careers



Upcoming Events



- State of .NET Online - .NET 8 – November 15th 2023
- <https://www.codemag.com/StateOfDotNet>



- CODE Presents – October 25th 2023:
Harnessing Large Language Models Locally
- <https://www.codemag.com/CODEPresents>



Free Subscription!

- The leading software development magazine written by expert developers for developers.
- All registered attendees will automatically receive a free digital subscription to CODE Magazine – no need to do anything, it'll happen auto-magically.
- Subscribers get CODE Focus issues free of charge!
- Please share this free subscription link:



<https://www.codemag.com/subscribe/cp10423>

Event Survey – Win \$100!

- Complete this very short 12 question survey for a chance at a \$100 Amazon Gift Card!
- Survey must be completed by 11:59pm ET on Friday 10/6/2023 to be eligible!
- THIS SLIDE WILL BE REPEATED AT THE END AND SURVEY LINK REPEATED IN THE CHAT WINDOW!

CODE Presents: Prompt Engineering - How to Talk to an AI Survey

The survey will take approximately 4 minutes to complete.

Thank you for attending! Please complete this brief 12 question survey to be eligible to win a \$100 Amazon Gift Card. Your survey must be completed by 11:59pm ET (UTC-4) on 09/02/2023 to be eligible to win! One entry per person please. Drawing will occur and the individual winner notified by 09/08/2023.

Thank you for attending! Please complete this brief survey. Yes, we still want to hear from you if you were unable to attend but watched the recording instead.

* Required

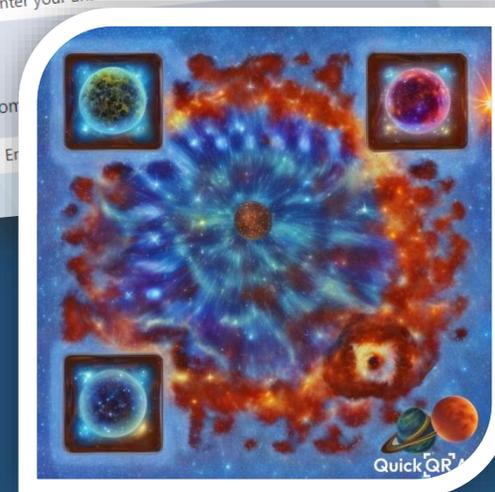
1. Full Name *

Enter your answer

2. Com

Er

<https://bit.ly/CP10423Survey>



Thank you for attending!



- Welcome new attendees!
- Welcome repeat attendees!
 - We now have quite a number of people who attended 50+ events
- Our AI presentations are the most popular we have ever done
 - ...and we have been at this for 30 years!
- Based on the attendee list, I am guessing we have a mix of developers, decision makers, and users

Agenda

1. Introduction

- Brief introduction to Generative AI and its importance
- Contextualizing Image Generation within Generative AI
- Goal of the presentation

2. Historical Evolution of Generative Models

- Brief history of generative models (with emphasis on image generation)
- Transition from early models to advanced architectures

3. Understanding Key Generative Models

- Introduction to Generative Adversarial Networks (GANs)
- Introduction to Variational Autoencoders (VAEs)

4. Running a Simple Generative Model

- Preparing the environment and data
- Training a simple model for image generation
- Discussing the parameters influencing the model's outcomes

5. Demonstration: Generating Images

- Demonstration of generating images using a pre-trained model
- Exploring the parameter settings and observing the outcomes

6. Discussion: Capabilities and Limitations

- Discussing the strengths and weaknesses of current generative models
- Future trends in generative AI for image generation

Introduction

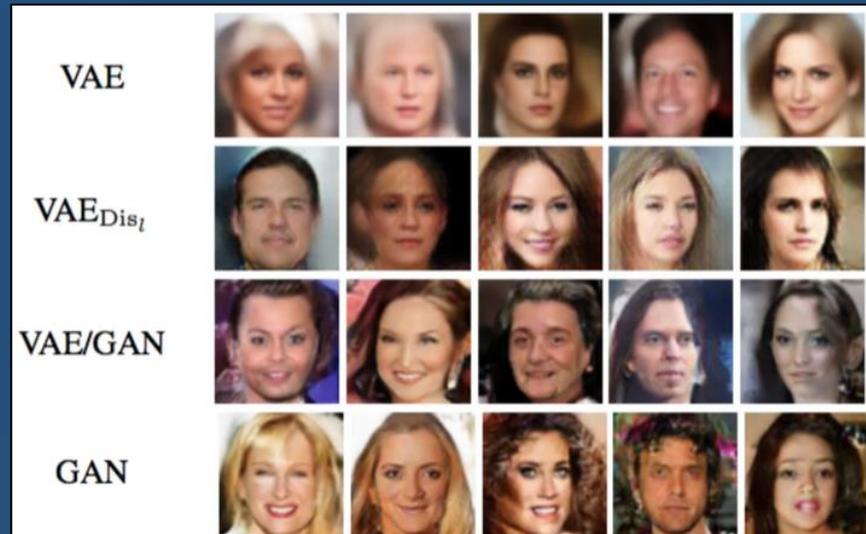
- Importance of Generative AI for Image Generation
- Contextualizing Image Generation
- Goal of the presentation

Historical Evolution

- Early Efforts: RBM (Restricted Boltzmann Machines)
- Introduction of Generative Adversarial Networks (GANs)
- Variational Autoencoders (VAEs)
- Improvement of GANs
- Flow based models
- Performance Benchmarks and App Driven Development
- Ethical Considerations
- Recent Advancements
- Open Source



GAN vs. VAE



Sik-Ho Tsang

Generative Adversarial Networks (GANs) **CODE** PRESENTS



Generative Adversarial Networks (GANs) **CODE** PRESENTS

- Generator – artist who tries to create a convincingly fake painting
- Discriminator – art critic who tries to tell if the painting is real or fake

For How Long?



Generative Adversarial Networks (GANs) **CODE** PRESENTS

- Learning Process: Cat and mouse game
 - Generator improves fakes
 - Discriminator improves at spotting fakes
- Training improves both sides
 - Feedback Loop
 - Convergence
 - Objective Measurement
 - Human Intervention (optional)



Variational Autoencoders (VAEs)

- More orderly and structured
- More like a skilled artist who studies and learns from real images to create new ones
- Study of images and their attributes (style, patterns, features, brushes, color palette)
- Create blueprint & recipe
- New Image Generation (batch)
- Evaluation



It's never just one!

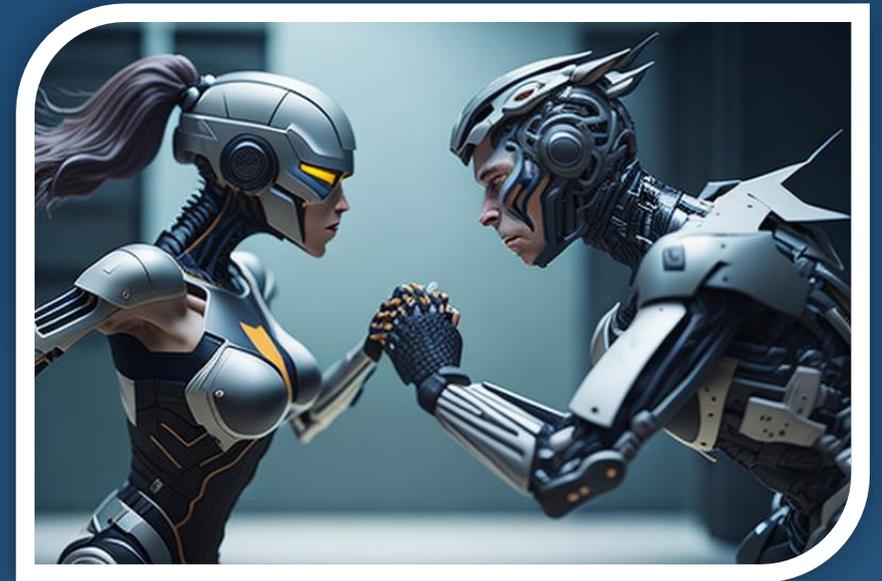
- All modern image generation services make use of both.
- Different emphasis and focus
- These concepts are **NOT** new!

AI Models and Timelines

Model	Notable AI products	Research Paper	Details
VAE	Amazon Alexa, Google Photos, Spotify	Dec 2013	Kingma & Welling "Auto-Encoding Variational Bayes"
GAN	AlphaGo, DALL-E, NVIDIA StyleGAN	Jun 2014	Goodfellow & Bengio et al. "Generative Adversarial Networks"
Flow models	Google Pixel 4, GPT-3, NVIDIA StyleGAN	May 2016	Dinh & Bengio et al. "Density estimation using Real NVP"
Transformer	BERT, Google Translate, OpenAI's GPT models	Jun 2017	Vaswani & Polosukhin et al. "Attention is All You Need"
Diffusion	DALL-E 2, Midjourney, Stable Diffusion	Jun 2018	Gregor & Weber et al. "Temporal Difference Variational Auto-Encoder"

Running a simple generative model

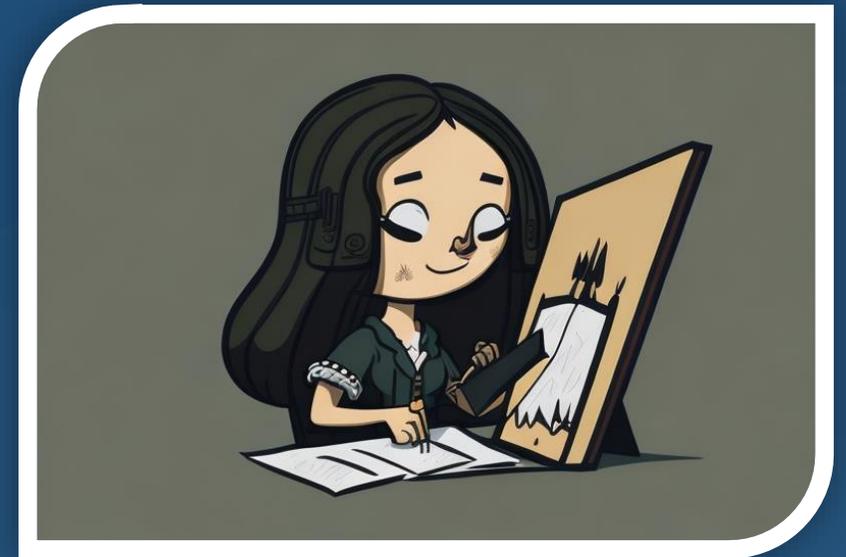
- Preparing the environment and data
 - Environment setup
 - Data collection
 - Data pre-processing
- Training a simple model for image generation
 - Model architecture
 - Training procedure
 - Evaluation during training



AI vs AI

Parameters influencing the model's outcomes

- Hyper-Parameters
- Regularization
- Loss Functions
- Optimization Algorithms & Algorithm Optimization
- Exploration of Latent Space

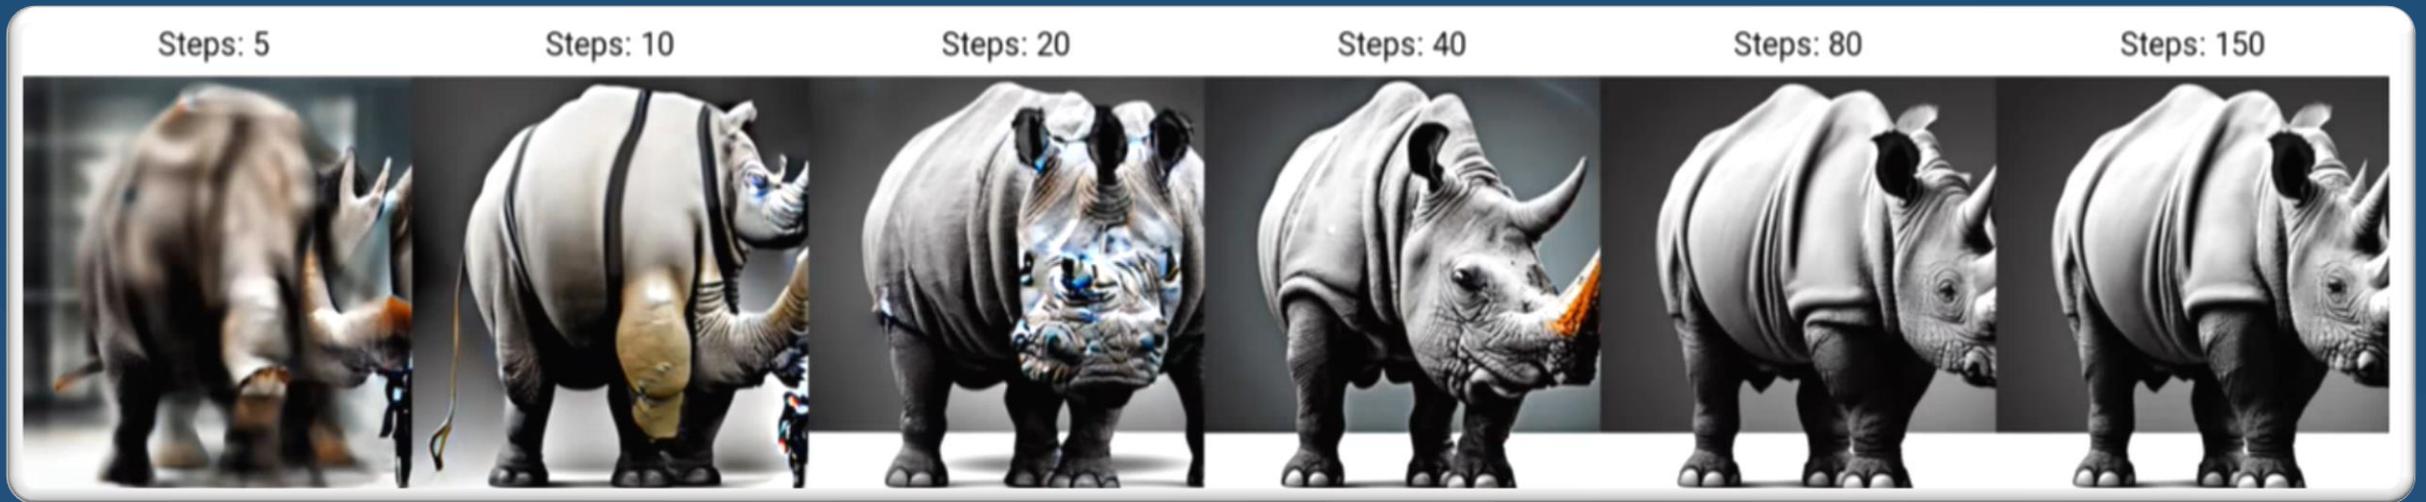


Generating images

- Pretrained Text Embeddings – Text To Vector
- Example: Rhino
 - Gray
 - 4 thick legs
 - 1 or 2 horns
 - Short tail
 - Skin texture
 - Postures
 - Etc.



Generating Images



Example results

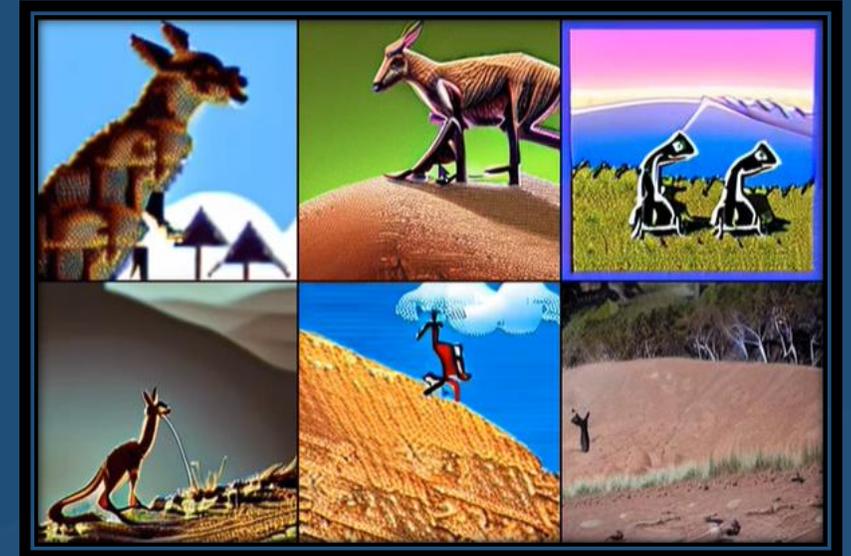
Blue car on moon



Red flower on ship



Kangaroo on mountain



6 images, each with 50 steps

Step Reduction



5 steps



50 steps

From noise to image



Prompt engineering

"astronaut cat, playing piano in space, black hole in the background, high quality, 4k render, dramatic light, photoreal, detailed, pov shot, volumetric lighting, golden hour"



"astronaut cat, playing piano in space, black hole in the background"



Leonardo.Ai with Alchemy

Service comparison

A dynamic action scene of a knight battling a dragon on a cliffside, with roaring waves below, inspired by the Romanticism movement, reminiscent of the works of Eugène Delacroix.



Midjourney



Nightcafe



Adobe Firefly



Leonardo

No settings changed, default parameters

Dreamshaper 8
(local)



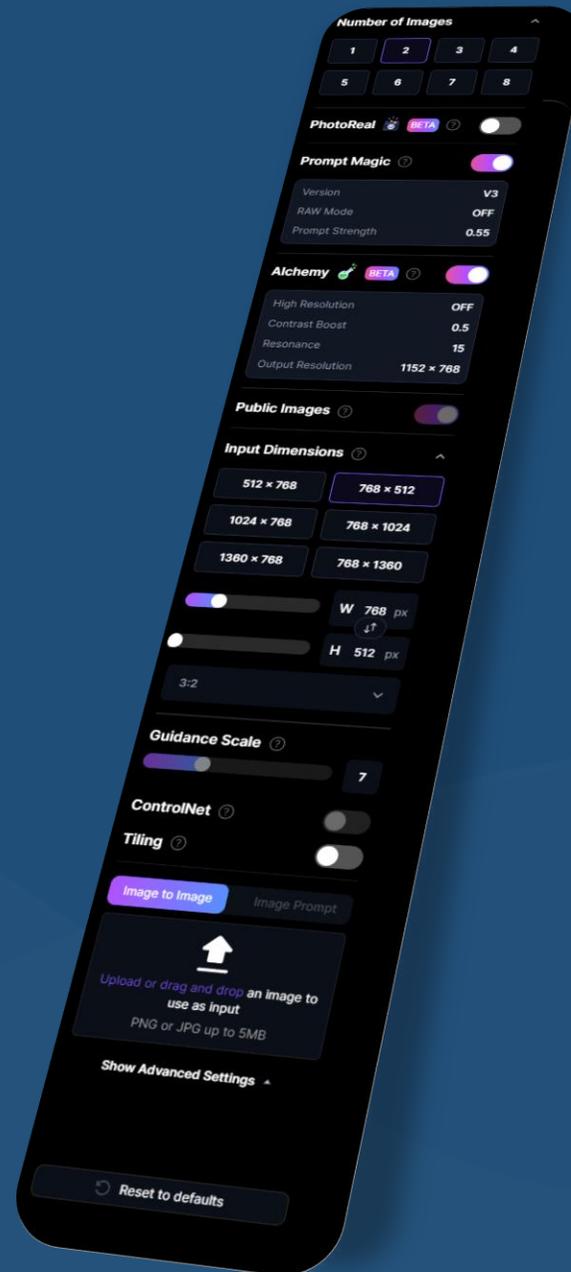
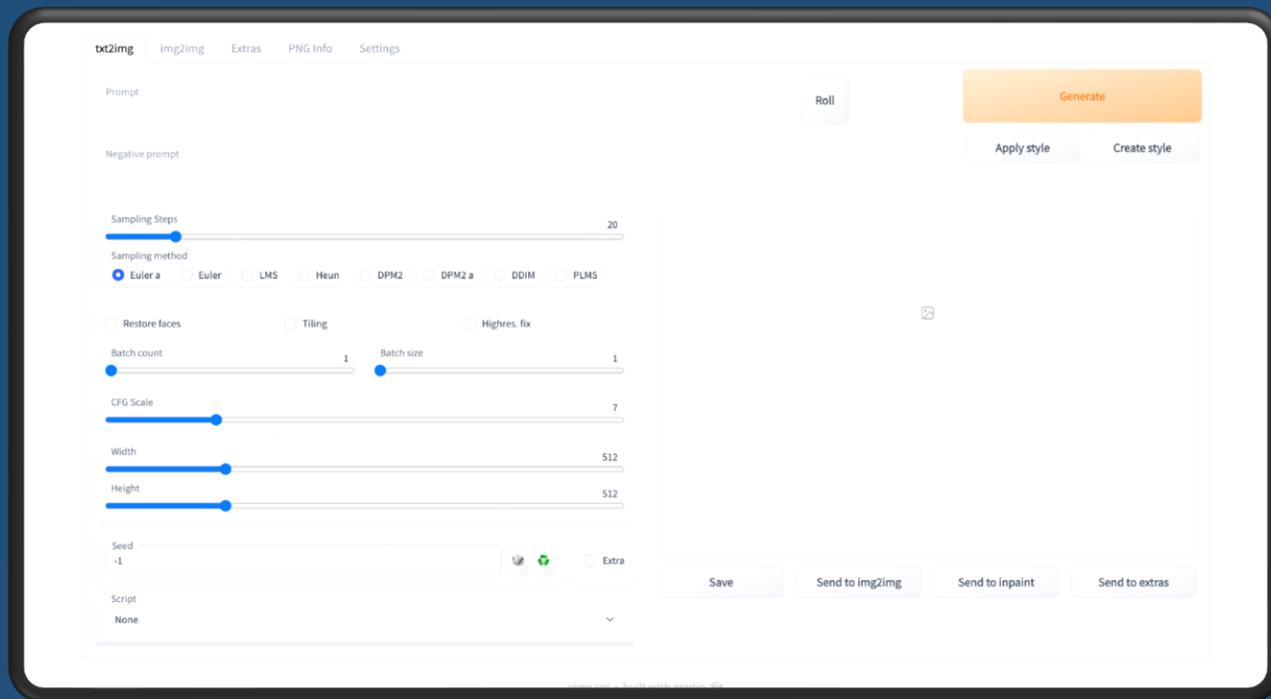
Local & personal cloud setup with Stable Diffusion

- For instance, on RunPod

The screenshot displays the RunPod GPU Cloud interface. On the left is a navigation sidebar with options like 'GPU Cloud', 'Serverless', 'Secure Cloud', 'Community Cloud', 'Browse Servers', 'Templates', 'MANAGE' (My Pods, My Storage, My Templates), 'Billing', 'Settings', and 'HELP' (FAQ, Contact). The main area shows a '2-3x faster than Ampere' banner and several GPU configuration cards. The top row features '8x H100 80GB' (\$2.3/hr), '1x RTX 8000 Ada' (\$1.39/hr), and '1x RTX 4090' (\$0.69/hr). Below this is a 'Previous Generation' section with configurations like '1x A100 80GB' (\$2.04/hr), '1x A100 SXM 80GB' (Unavailable), '1x A40' (\$0.79/hr), '1x A6000' (\$0.79/hr), '1x RTX 3090' (\$0.44/hr), '1x RTX A5000' (\$0.44/hr), '1x RTX A4500' (\$0.36/hr), and '1x RTX A4000' (\$0.34/hr). Each card lists specifications like VRAM, RAM, and vCPU, along with pricing and availability status.

Secure Cloud		1x GPU \$/hr	8x GPU \$/hr
A100 80GB 80 GB VRAM 125 GB RAM 8 vCPU	Spot	-	-
	On-Demand	1.99	15.92
	1 Month	-	-
	3 Month	-	-
A100 SXM 80GB 80 GB VRAM 125 GB RAM 16 vCPU	Spot	-	-
	On-Demand	2.19	17.52
	1 Month	-	-
	3 Month	-	-
H100 80GB SXM5 80 GB VRAM GB RAM vCPU	Spot	-	-
	On-Demand	4.49	35.92
	1 Month	-	-
	3 Month	-	-
H100 80GB PCIe 80 GB VRAM 125 GB RAM 16 vCPU	Spot	-	-
	On-Demand	4.29	34.32
	1 Month	-	-
	3 Month	-	-
A40 48 GB VRAM 50 GB RAM 4 vCPU	Spot	-	-
	On-Demand	0.79	6.32
	1 Month	-	-
	3 Month	-	-
L40 48 GB VRAM 250 GB RAM 16 vCPU	Spot	0.69	5.52
	On-Demand	1.14	9.12
	1 Month	1.08	8.64
	3 Month	1.02	8.16
RTX 6000 Ada 48 GB VRAM 61 GB RAM 8 vCPU	Spot	-	-
	On-Demand	1.14	9.12
	1 Month	-	-
	3 Month	-	-
RTX A6000 48 GB VRAM 50 GB RAM 4 vCPU	Spot	0.49	3.92
	On-Demand	0.79	6.32
	1 Month	0.75	6
	3 Month	0.71	5.68
RTX 3090 24 GB VRAM 30 GB RAM 4 vCPU	Spot	-	-
	On-Demand	0.44	3.52
	1 Month	-	-
	3 Month	-	-
RTX 4090 24 GB VRAM 29 GB RAM 4 vCPU	Spot	0.49	3.92
	On-Demand	0.74	5.92
	1 Month	-	-
	3 Month	-	-

Web UIs



Hugging Face

The screenshot shows the Hugging Face Spaces page. At the top, there is a navigation bar with the Hugging Face logo, a search bar for models, datasets, and users, and links for Models, Datasets, Spaces, Docs, Solutions, Pricing, Log In, and Sign Up. Below the navigation bar, the 'Spaces' section is highlighted, with a sub-header 'Discover amazing ML apps made by the community!' and a 'Create new Space' button. A search bar for Spaces and a 'Sort: Trending' dropdown are also present. The main content area is divided into three sections: 'Spaces of the week', '# All running apps, trending first', and a grid of application cards. Each card displays the app name, a heart icon with a count, the hardware it runs on, the creator's name, and the time since it was updated.

App Name	Hardware	Creator	Time Ago	Heart Count
LoRA Roulette	Running on ZERO	multimodalart	1 day ago	91
Blind Chat	Running on T4	mithril-security	about 4 hours ago	73
Upside-Down-Diffusion	Running on T4	AP123	about 14 hours ago	104
DeciDiffusion-v1-0	Running on A10G	Deci	16 days ago	66
Nougat Transformers	Running on T4	hf-vision	3 days ago	50
Stable Diffusion 2-1	Running on CPU UPGRADE	stabilityai	14 days ago	9.11k
Explore Clinical & Biomedical Language Models	Running on T4	hf4h	15 days ago	77
Lilac	Running on CPU UPGRADE	lilacai	2 days ago	24
IllusionDiffusion	Running on A10G	AP123	about 12 hours ago	1.23k
AI Comic Factory	Running on CPU UPGRADE	jbilcke-hf	about 6 hours ago	1.83k
Shopify Image Background Replacement	Running on A100	Shopify	about 2 hours ago	132
Open LLM Leaderboard	Running on T4	HuggingFaceH4	1 day ago	5.45k

Other Announcements

AI Consulting Services



CODE CONSULTING Home Magazine Consulting Staffing Training Framework VFP Conversion Login

Executive Briefing (AI)

Are you wondering how AI can help your business? Do you worry about privacy or regulatory issues stopping you from using AI to its fullest? We have the answers! Our Executive Briefings provide guidance and concrete advice that help decision makers move forward in this rapidly changing Age of Artificial Intelligence and Copilots!

We will send an expert to your office to meet with you. You will receive:

1. An overview presentation of the current state of Artificial Intelligence.
2. How to use AI in your business while ensuring privacy of your and your clients' information.
3. A sample application built on your own HR documents - allowing your employees to query those documents in English and cutting down the number of questions that you and your HR group have to answer.
4. A roadmap for future use of AI catered to what you do.

A Briefing Combined with Solid Deliverables!

We will send an expert to your office to meet with you. You will receive:

1. An overview presentation of the current state of Artificial Intelligence.
2. How to use AI in your business while ensuring the privacy of your data.
3. A sample application built on your own HR documents - allowing your employees to query those documents in English and cutting down the number of questions that you and your HR group have to answer.
4. A roadmap for future use of AI catered to what you do.

Executive Briefings are set up specifically for you and usually take several hours and up to a day. It is up to you whether you want to include only C-level Executives or your entire company. The cost is **\$5,000** (which is a discounted special rate, compared to the normal price of \$7,500) plus expenses if we have to travel to your location, or stay over night. This includes the concrete deliverable of a first implementation of AI for your business (see #3 above, although we may also agree that something slightly different is more appropriate for your organization).

Contact us today to set up your individual Executive Briefing!

Setting Directions for the Future of Your Organization

We are experiencing a time of extreme change in technology. While one might argue that this has always been true, it is clear that current developments have a drastic impact on all organizations at a scale that is orders of magnitude beyond what we have experienced in years past.

This represents a time of extreme opportunity, as well as a time of considerable risks!

Artificial Intelligence (AI) continues to revolutionize industries across the globe, offering immense potential for organizations to enhance operational efficiency, drive innovation, and achieve competitive advantage. Conversely, ignoring these technological trends means that even very recent technology is quickly turning into "legacy" and your competitors now have the opportunity to leapfrog you. Especially the advent of "The Age of Copilots" (powered largely by the onset of Large Language Models - LLMs) promises a productivity improvement of upwards of 50% in some areas year after year, and that isn't the only trend to pay attention to.

What does all of this mean specifically for your organization?

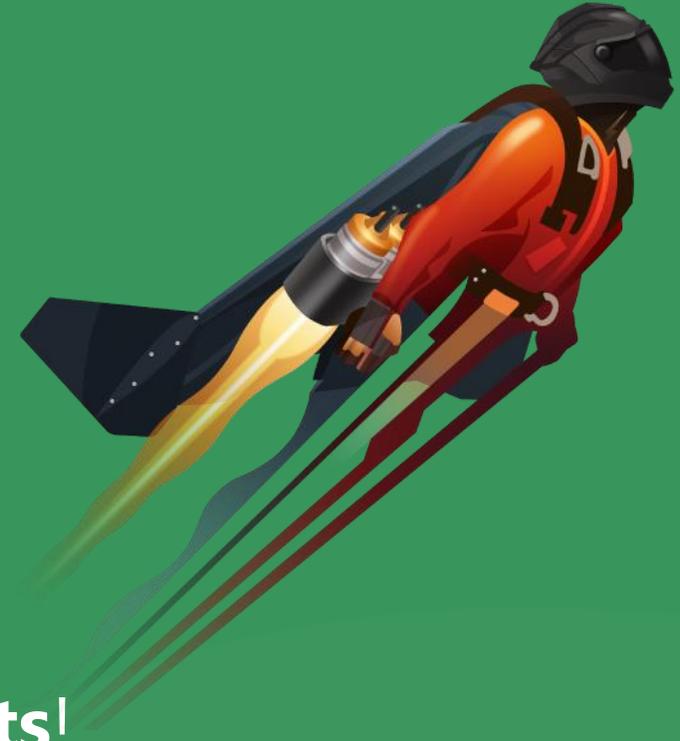
- Check out our new **Executive Briefing Offer!**
- We can help with your AI needs
- What does AI mean for you?
- "Skunk Works" Projects



- www.codemag.com/AI
- www.codemag.com/ExecutiveBriefing

CODE Staffing

- Disrupting the world of staffing!
- Giving our customers the ability to have staff on par with Silicon Valley companies...
- ...and our employees a work environment in a bleeding-edge tech company with the industry leading benefits!
- www.codestaffing.com



We Need You!

- We are currently “hiring on all cylinders” across the entire CODE Group Organization!
- Check out www.codestaffing.com/careers



Upcoming CODE Training Classes



- Public AI Training Classes
 - Introduction to Development with Open AI, ChatGPT, Azure Open AI, and more...
 - Building Copilots for your own apps
- Classes are also available as individualized private classes and mentoring sessions



<https://www.codemag.com/training>

Upcoming Events



- State of .NET Online - .NET 8 – November
- <https://www.codemag.com/StateOfDotNet>

- CODE Presents – Monthly – Mostly AI
- <https://www.codemag.com/CODEPresents>



Recordings & Slide Decks

- State of .NET Webinar Series
 - codemag.com/StateOfDotNet
- CODE Presents Webinar Series
 - codemag.com/CODEPresents

Recent Event Recordings

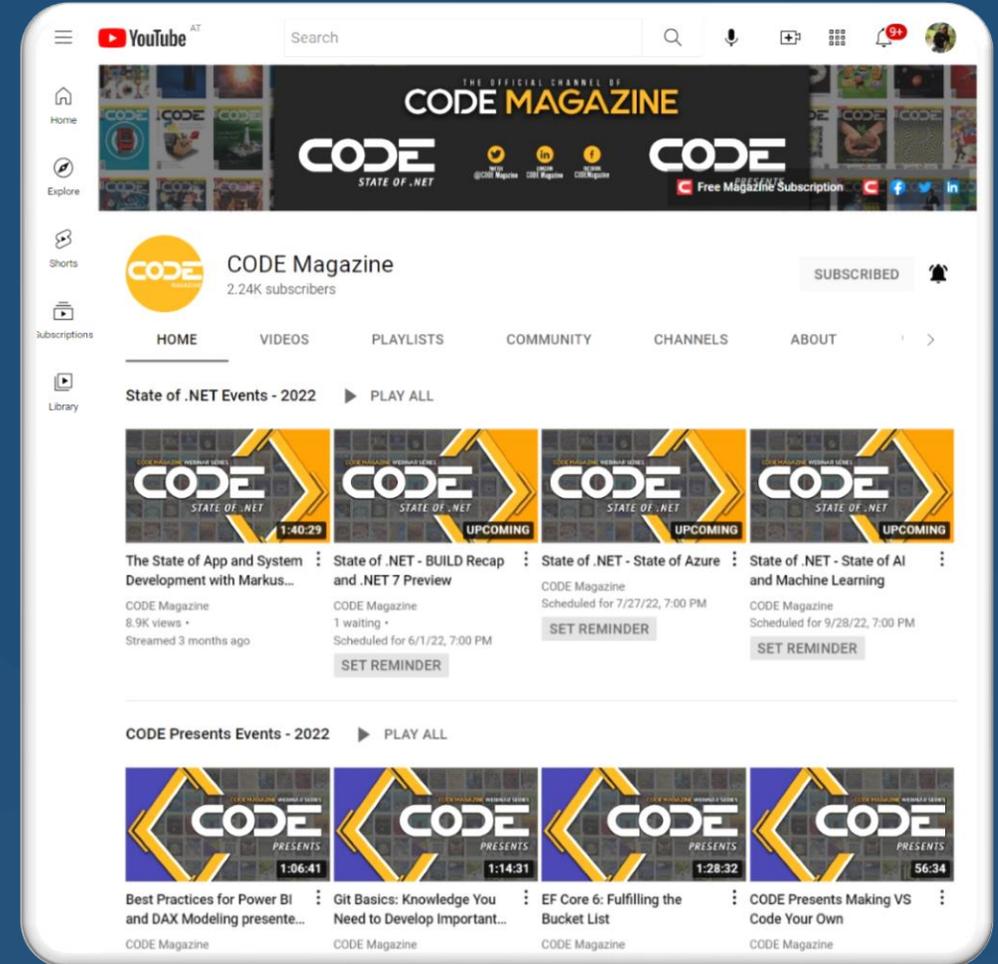
<p>State of .NET - The State of .NET App and System Development</p> <p>2/23/2022</p>  <p>Thumbnail for the recording "State of .NET - The State of .NET App and System Development" dated 2/23/2022. The thumbnail features the CODE logo and the text "STATE OF .NET" on a background of repeating CODE logos.</p>	<p>State of .NET - .NET Conf Recap</p> <p>11/17/2021</p>  <p>Thumbnail for the recording "State of .NET - .NET Conf Recap" dated 11/17/2021. The thumbnail features the CODE logo and the text "STATE OF .NET" on a background of repeating CODE logos.</p>
<p>State of .NET - .NET 6 Preview</p> <p>10/27/2021</p>  <p>Thumbnail for the recording "State of .NET - .NET 6 Preview" dated 10/27/2021. The thumbnail features the CODE logo and the text "STATE OF .NET" on a background of repeating CODE logos.</p>	<p>State of .NET - Azure Architecture</p> <p>9/29/2021</p>  <p>Thumbnail for the recording "State of .NET - Azure Architecture" dated 9/29/2021. The thumbnail features the CODE logo and the text "STATE OF .NET" on a background of repeating CODE logos.</p>

Subscribe to our YouTube Channel!



- Subscribe and “ring that bell” to never miss any of our content!

- www.youtube.com/Codemag



Event Survey – Win \$100!

- Complete this very short 12 question survey for a chance at a \$100 Amazon Gift Card!
- Survey must be completed by 11:59pm ET on Friday 9/8/2023 to be eligible!
- THIS SLIDE WILL BE REPEATED AT THE END AND SURVEY LINK REPEATED IN THE CHAT WINDOW!

CODE Presents: Prompt Engineering - How to Talk to an AI Survey

The survey will take approximately 4 minutes to complete.

Thank you for attending! Please complete this brief 12 question survey to be eligible to win a \$100 Amazon Gift Card. Your survey must be completed by 11:59pm ET (UTC-4) on 09/02/2023 to be eligible to win! One entry per person please. Drawing will occur and the individual winner notified by 09/08/2023.

Thank you for attending! Please complete this brief survey. Yes, we still want to hear from you if you were unable to attend but watched the recording instead.

* Required

1. Full Name *

Enter your answer

2. Con

Er



<https://bit.ly/CP10423Survey>

Q&A

Contact us with questions!

CODE/EPS Contact:

www.codegroup.io

info@codemag.com

facebook.com/codemag

twitter.com/codemagazine

Presenter Contact:

ottod@codemag.com