# How AI is Enabling a Creativity Renaissance

Ben Falchuk, Ph.D.

Peraton Labs, NJ, USA

Peraton

Email: bfalchuk@peratonlabs.com | ben.falchuk@gmail.com

Presented at IARIA ACHI'21

# Author Bio

**Ben Falchuk** is a Senior Scientist at *Peraton Labs (formerly Perspecta)*. He holds a **Ph.D.** (Computer Engineering) from the University of Ottawa, an **M.Sc**. (Computer Science) from Carleton University, and a **B.Math**. (Applied Mathematics) from the University of Waterloo, Canada.

His work focuses largely on human computer interaction techniques, design, and software. Dr. Falchuk is active in the research community as author and program committee member. He also builds creative web experiments and provides technical expertise to nonprofit organizations.

*Peraton Labs* provides software and advanced R&D to customers in realms including cybersecurity, electronic warfare, mobility and networking.



Peraton

# Outline

- Motivation
- Background
- Scope
- Survey*
- Trends
- Conclusions

* This presentation only summarizes; readers should **refer to the accompanying paper**.



Peraton

# Motivation

- Advances in computer software and hardware are resulting in innovative products that support the *creative class** in increasingly effective ways:
  - E.g., products from Adobe, Autodesk, Microsoft, as well as open source products like Blender, and apps such as PicsArt, gimp, and Procreate, to name a few
- Deep learning and AI have seen incredible gains, relatively recently
  - Spurred by improvements in hardware (GPU), software, open data sets, and more
- As AI and creative tools come together, we find ourselves at a moment of explosive growth (in the power, capabilities, and diversity of hardware and software applications)
- Will AI-backed software turn us all into creative, multi-talented, Leonardo da Vinci's?
  - What are recent advances and products?
  - In which fields do we see AI making in-roads?
  - How will humans and AI interwork in these fields?

* As defined by R.Florida, a sort of socioeconomic class comprised of occupations such as science, engineering, computer programming, arts, and design

# Background

- The following terms are used often in this survey:

- Neural network – a nested mathematical function with a particular form, modeled loosely upon constructs of the human brain
- Deep learning – the utilization of Artificial Neural Networks (ANN) that have more than one layer between input and output layers
- Generative Adversarial Network (GAN) – a class of neural networks often implemented in pairs, each contesting with each other in a sort of zero-sum game
- Artificial Intelligence – a branch of computer science dealing with the simulation of intelligent behavior in computers [via Miriam Webster]
- Creativity – the production and/or use of original and or novel/unexpected ideas that have a distinct effect on human viewers [via Cambridge dict.]
  - AI may be a ***source of*** creativity, or may be an ***inspiration for*** (e.g., muse) human creativity; this distinction is quite immaterial here because, either way, it is the result that is of interest



Peraton

# Scope

- In the remainder of this presentation we survey the following realms:

- Writing / Poetry
- Cuisine
- Painting/drawing
- Photography/portraiture
- Music
- AI Platforms

- We list and describe the capabilities of several notable technologies in each category*
  - We focus on interesting and novel human-AI interactions
  - The list is not exhaustive, rather it selects exemplary and state-of-the-art systems from each domain
  - **This presentation only summarizes; refer to the accompanying paper for further detail**

# Drawing and Painting

- Algorithmic style-transfer (via GAN) is proliferating
  - "Draw some dinosaurs made of flowers",
  - "Draw this scene in the style of Vincent van Gogh"
- Digital paintings can be "assembled" by algorithms (generative art has a long history)
- AI can fill in missing details of images (e.g., in landscapes, and human faces) based on patterns gleaned from large training sets
  - A useful capability in many realms, including creative ones
- AI tools can stream out new artworks "in the style of Z" inexhaustibly and on a continual basis
- Key Human-AI aspects:
  - AI is tireless and can create sources of **inspiration** for human artists
  - Human users can rough-in ideas, while AI can work on filling-in tedious or tiresome details

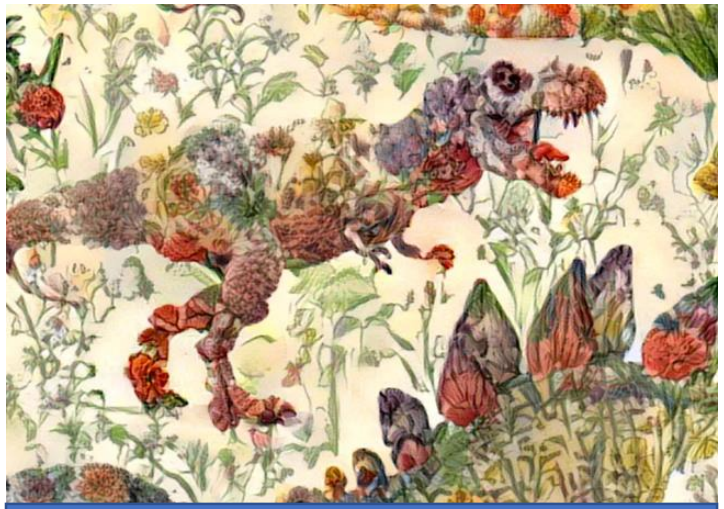

Dinosaurs of flowers, C.Rodley

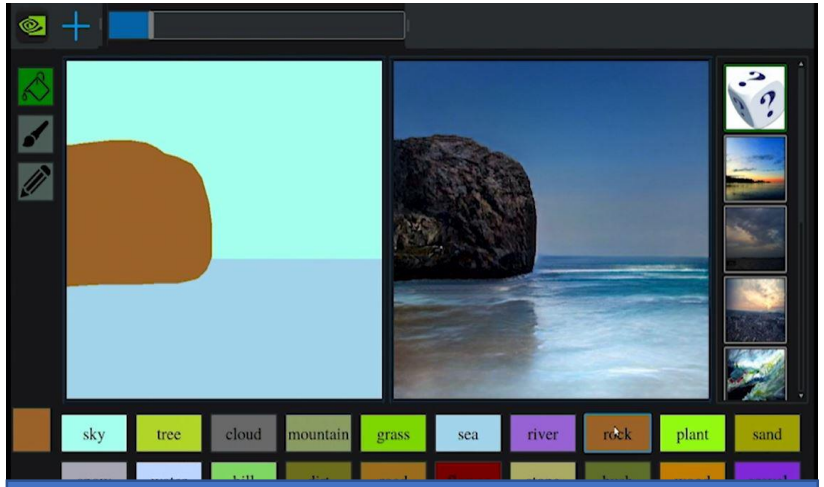


Human input (left), GauGAN output (right)



Peraton

# Drawing and Painting

"La Famille d Belamy" AI artwork sold for over $400,000 in auction (in 2018)

DeepDream – in this image a network trained to perceive dogs iterates over Mona Lisa



Peraton

# Writing and Poetry

- A Generative Pre-Training (GPT) model is a (language or image) model trained on real text (or images) and able to generate convincing output. OpenAI has released several improving variants of GPT
- USC/Univ.Washington work (funded by DARPA) results in ability to create a poem about a given topic.
- AI-generated college essays resulted in mostly passing grades (except creative writing) when graded by experts
- Verse by Verse is trained on the corpus of scores of American authors (>25 millions words), allows users to collaboratively create a new poem in the style of selected authors
- DALL-E by OpenAI can generate images from inputted text by leveraging GPT-3 and Image GPT
- Key Human-AI aspects:
  - These tools can (and do) enable new modes of human-AI interaction for literature. E.g., collaborative creation of text, interactive illustration of human words, adding a degree of rhyme (poetic stylings), and more

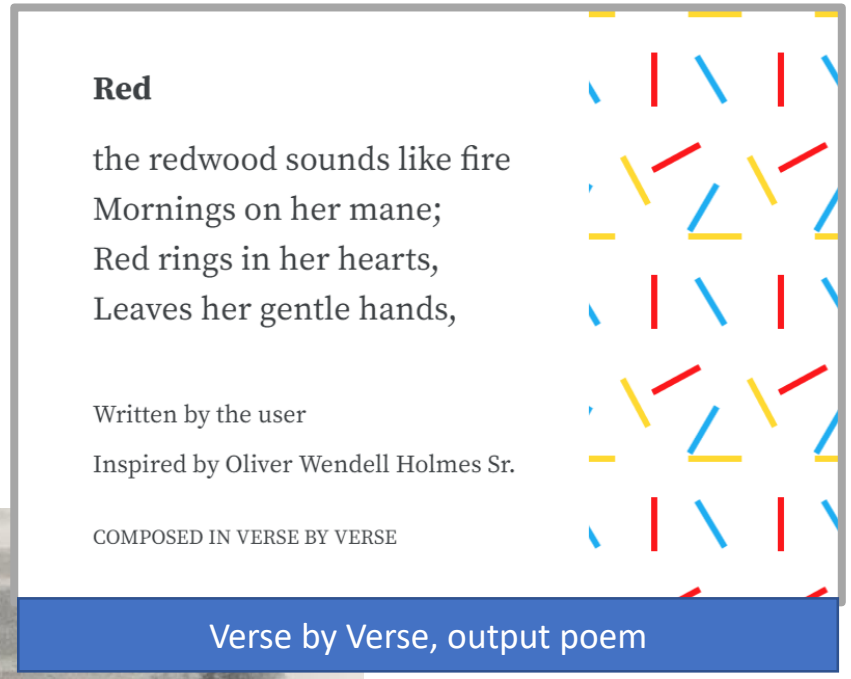


Verse by Verse, output poem

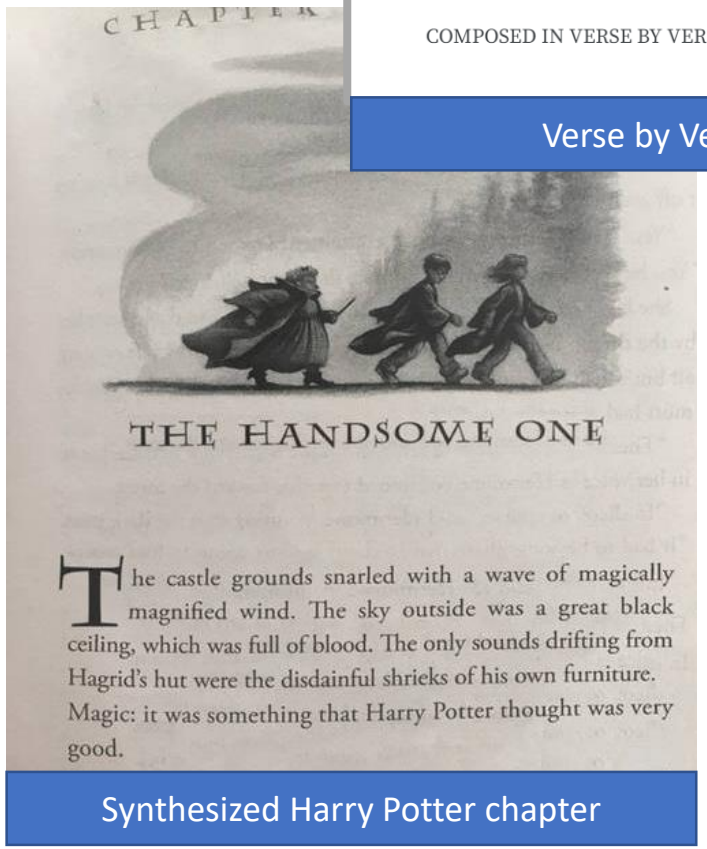


Synthesized Harry Potter chapter



Peraton

# Writing and Poetry

DALL-E outputs when prompted for “armchair in the shape of avocado”


Please feed the Lions exhibit, generative poetry, London, E.Devlin



Peraton

# Culinary Arts

- Experiments with GPT-3 result in reasonable-sounding – but impractical – recipes
  - i.e., a recipe is more than a string of interesting phrases; it must "work" on culinary/practical levels
- The KitcheNette app finds novel food pairings thanks to large training sets of existing recipes
- IBM Watson supercomputer – trained on recipe corpus – has created AI food recipes that have some real mass appeal
- Image-to-recipe (based on neural networks trained on food, recipes and topic networks) can identify ingredients from only a photo of food on a dish
- Key Human-AI aspects:
  - AI can almost certainly be an effective sous-chef, recommending pairings and ingredients
  - AI tools may recognize what a chef is doing (image-to-recipe) and offer creative variants on the recipe from a vast training set.

AI-generated recipe called, "Caymanian Plantain Dessert"

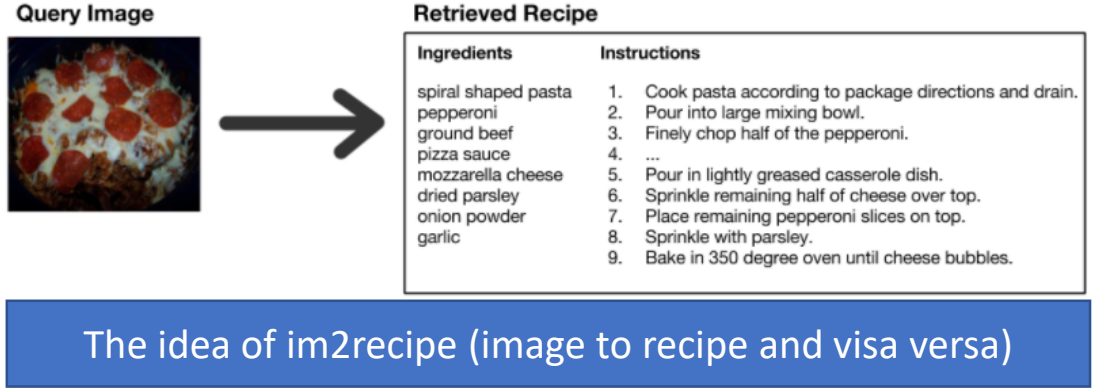


The idea of im2recipe (image to recipe and visa versa)



Peraton

# Photography, Portraiture

- AI is supporting illustrative works with new abilities to upscale images
- GANPaint and similar projects can convincingly replace elements in photos with other ones
- StyleGAN2 generator creates images of human faces while controlling selected style aspects like hair, nose, eyes
  - E.g., Toonify
- Synthetic faces can be created with photographic realism
  - So-called deepfakes are potentially dangerous
- Static portraits can be make to move, talk and/or walk, generally by leveraging large training sets and applying contextually relevant attributes
- Key Human-AI aspects:
  - AI is already prevalent in photography – in particular as related to human face and body (for which high-quality training data is available)
  - The AI can become a sort of "muse" to the human artist, replacing and augmenting images in iterative fashion


GANPaint – replaces a tower with a tree, maintaining semantics of image

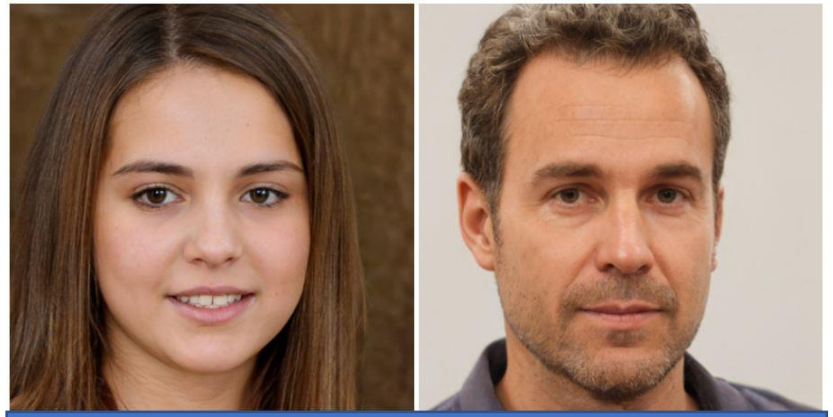

Exemplary faces generated by an AI in the tool, *This Person does not Exist*



Peraton

# Photography, Portraiture

*Toonify* is an application of StyleGAN that transfers cartoon styling to photographs

Artbreeder outputs (each is “bred” from two previous (unseen) inputs)



Peraton

# Music

- OpenAI's Jukebox was training on >1Million songs, lyrics across genres, can predict audio tokens and create convincing new songs
- DeepSinger was trained specifically on singing and voice data and can synthesize new singing voices in multiple languages
- DrumBot and Air Guitar webapp experiments show that AI-based collaboration and (computer-vision) gesture-recognition are new – and effective - underpinnings
- Google's Magenta can generate new songs in the style of other trained sample and recently created a convincing new Nirvana song (along with AI-generated lyrics)
- Visualization AI's – such as Deep Music Visualizer - are popular on demo-scenes wherein lucid visual accompaniments are made to songs
- Google's FreddieMeter AI judges how close a user's vocals are to those of Freddie Mercury across a number of parameters
- Key Human-AI aspects:
  - AI-composers may support humans with riffs, and bridges within songs
  - AI generated music has been hit and miss, sometimes lacking the musical cohesion one expects

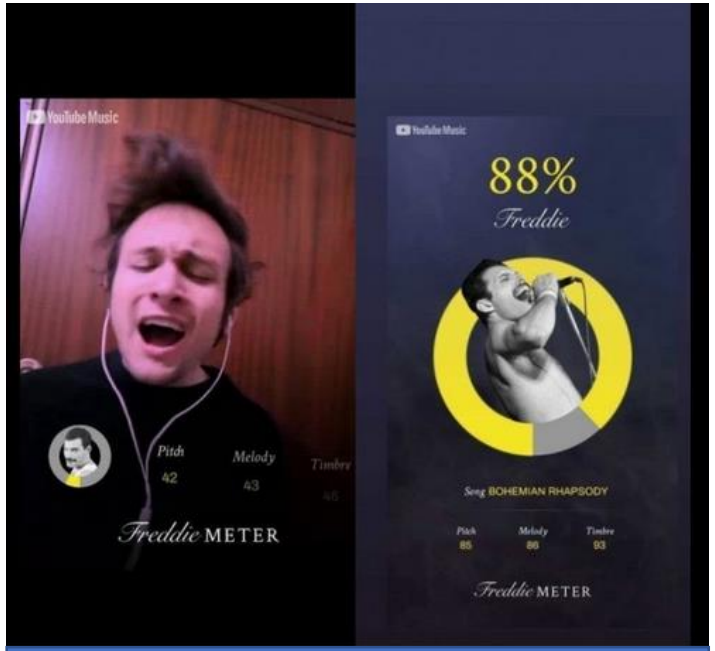


A singing user (left) being ranked by FreddieMeter webapp (right)

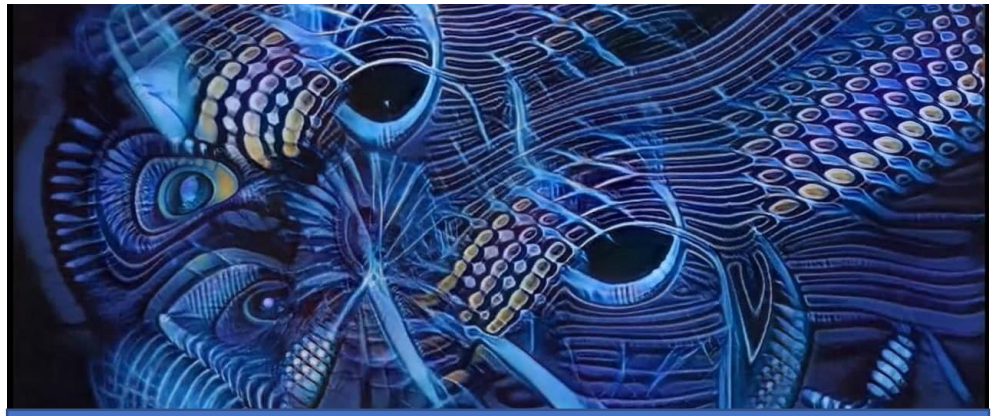

Exemplary image from Lucid Sonic Dreams



Peraton

# Music

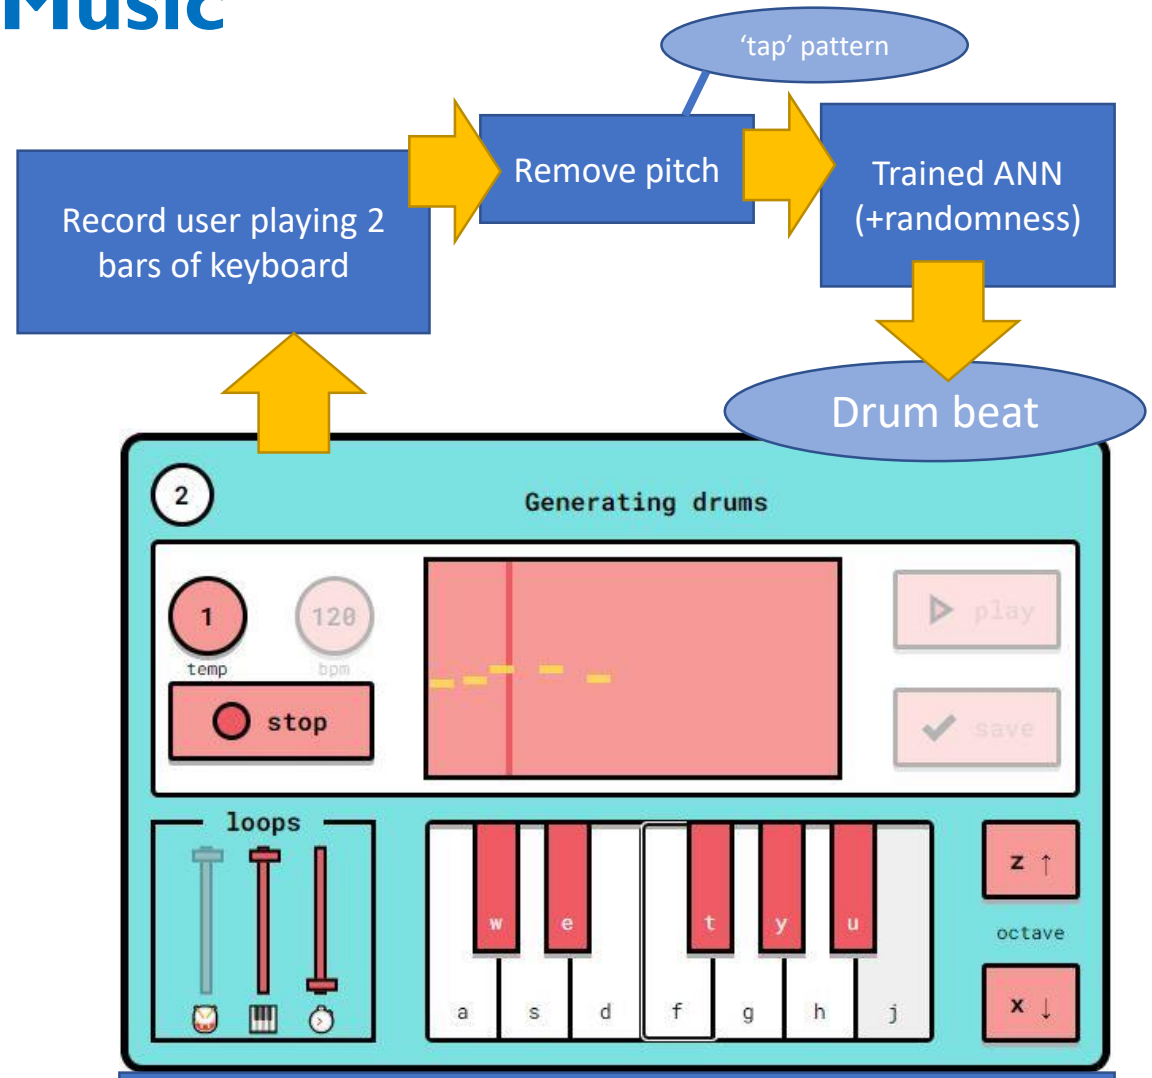


Drumbot uses AI to drum along with a musical pattern created by human user – the ANN creates drum beats from 'tap' patterns

*"Air guitar"* codepen demo that uses pose recognition (with browser webcam) to create music



Peraton

# AI Platforms

- To encourage buy-in within tech-ecosystem and/or to showcase the capabilities and scope of their prowess, companies are creating AI platforms for creative users
  - While coding is still a necessary aspect, more and more code-free AI seems to be a trend
    - Lowers the bar for adoption and enables a feedback loop when more creative users employ the tools
  - Abilities like drag-and-drop (of datasets) and auto-visualization (of neural networks) greatly assist creative developers who do not have a background in data science or machine learning
- Examples*:
  - Lobe (owned by Microsoft)
  - SageMaker (owned by Amazon)
  - Runway ML
  - MediaPipe
  - TensorFlow
- Key Human-AI aspects:
  - Becoming much easier to use, for all levels of coding capabilities

* Readers should refer to figures and references in the accompanying paper.

# Trends

- We see the following trends:

1. Improving access to massive data sets for training, and/or pre-trained models
2. AI techniques and algorithms are faster and more accessible than ever before
   - Nearly "code free" platforms and toolkits (e.g., Lobe) are on the rise
3. Ever-increasing potential for human-AI collaborative flows in which each actor contributes unique value
4. Human-AI interactions have ever-improving accessibility
   - Via UI and web technology break-throughs, even in browsers, and in mobile devices; voice and NLP remain a highly viable means for non-UI conversational interaction
5. Single-domain AI tools currently dominate the scene, while those with domain-crossing skills are more rare

# Conclusions and Future Work

- AI is being applied in realms of human creativity with fascinating results
- AI can be a direct partner of - or more indirect source (e.g., a muse) for - human creativity
- While artists have been using AI for decades, the proliferation of massive open data sets, efficient frameworks for transforming and ingesting data, availability of pre-trained models (e.g., PoseNet) is newer
  - Building and deploying AI apps from ground-up is also becoming simpler
- But challenges remain with regards to AI and Human-AI systems:
  - Ensuring data sets are fair and free from bias
  - Understanding the roles of consciousness, personality, ethics (and more) with regards to human-AI collaboration
  - Building federations of AI tools that can compute across several domains at once (e.g., both visual and textual intelligence)
- Our work at Peraton Labs work focuses on AI and deep learning fundamentals, human-interfaces, and the application of AI across domains such as cybersecurity, electronic warfare, mobility and networking

# How AI is Enabling a Creativity Renaissance

## Thank You!  Questions?

Ben Falchuk
Contact emails: bfalchuk@peratonlabs.com | ben.falchuk@gmail.com



IARIA

Peraton