Modest, artistic, and radical solutions to the environmental impact of image-generating machine learning

Kehui Li, Jess MacCormack, and Laura U. Marks
Simon Fraser University

Computing within Limits 2026, June 25, 2026

Background
Calculating the Carbon Footprint of Streaming Video (2020-2022) Marks, Stephan Makonin, Alejandro Rodriguez-Silva, Radek Przedpełski

The Small File Media Festival, since 2020
https://smallfile.ca –a tiny website designed by SFMF's Joni Schinkel
Max. bitrate 1.44 MB/minute
Festival live at The Cinematheque, Vancouver, and streaming at low impact worldwide

International workshops: Beirut, Cairo, Dhaka, Dublin, Kashmir, Las Puntas, Mexico City, Tehran

The Machine Learning Environmental Impact research team, Simon Fraser University

A subgroup of Smallfile Generative Art project, led by robotics sound composer Arne Eigenfeldt, to create generative media art inspired by small-file media

Laura U. Marks, media scholar, PI of Tackling the Carbon Footprint of Streaming Media
Stephen Makonin, computer scientist, founder of Computational Sustainability Lab
Jess MacCormack, artist, PhD student in Contemporary Arts
Kehui Li, MA student, Computer Science

Sharing approaches and methods across disciplines

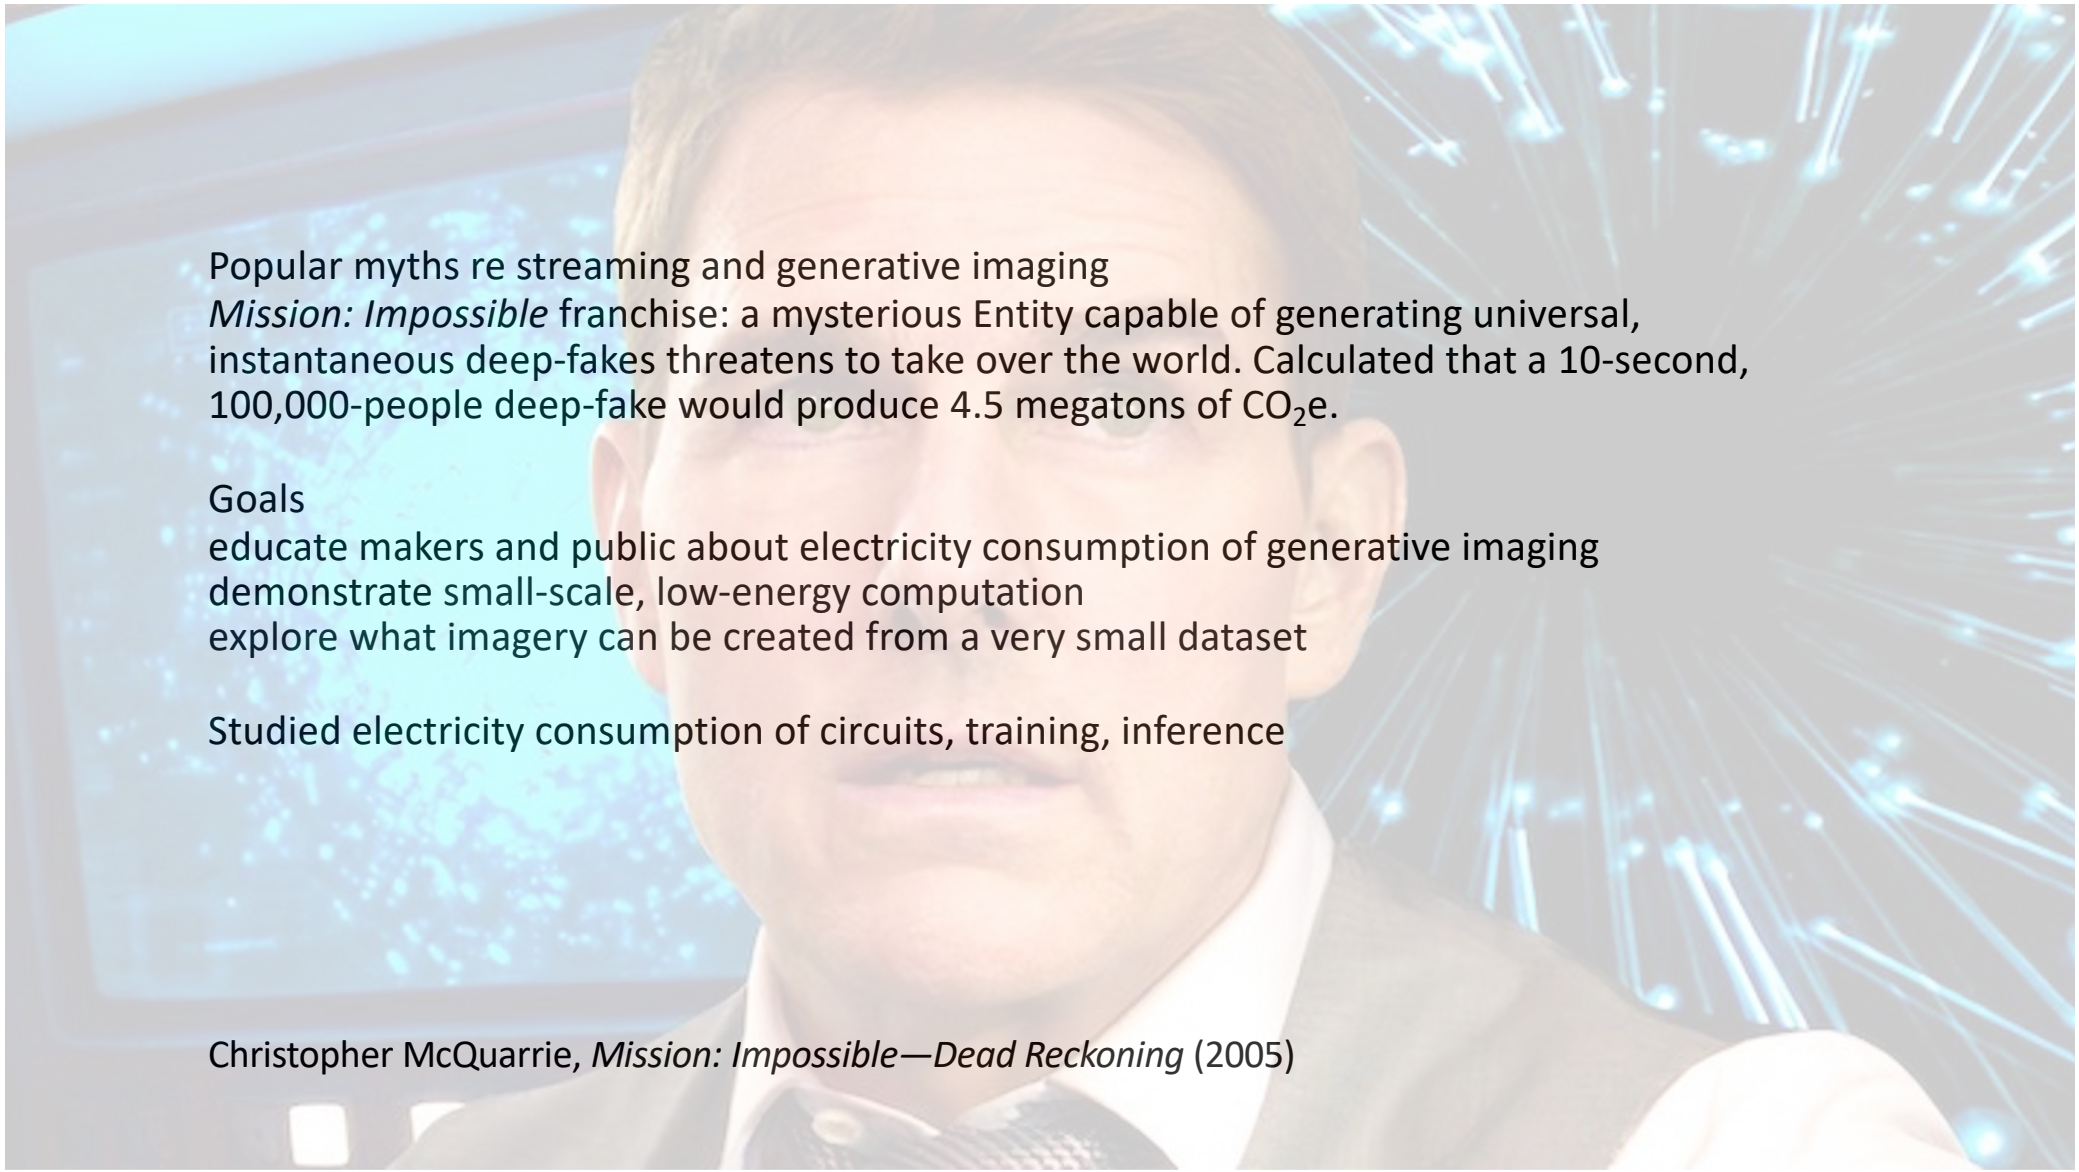

Popular myths re streaming and generative imaging
Mission: Impossible franchise: a mysterious Entity capable of generating universal, instantaneous deep-fakes threatens to take over the world. Calculated that a 10-second, 100,000-people deep-fake would produce 4.5 megatons of CO_2e.

Goals
educate makers and public about electricity consumption of generative imaging
demonstrate small-scale, low-energy computation
explore what imagery can be created from a very small dataset

Studied electricity consumption of circuits, training, inference

Christopher McQuarrie, *Mission: Impossible—Dead Reckoning* (2005)

We’re ambivalent about introducing methods to create a tiny image generator.
Rebound effects of “efficient” models are too great.
The best generative imaging is no generative imaging.

However, we are proud of our work to develop this speculative model, which is also a conceptual artwork.

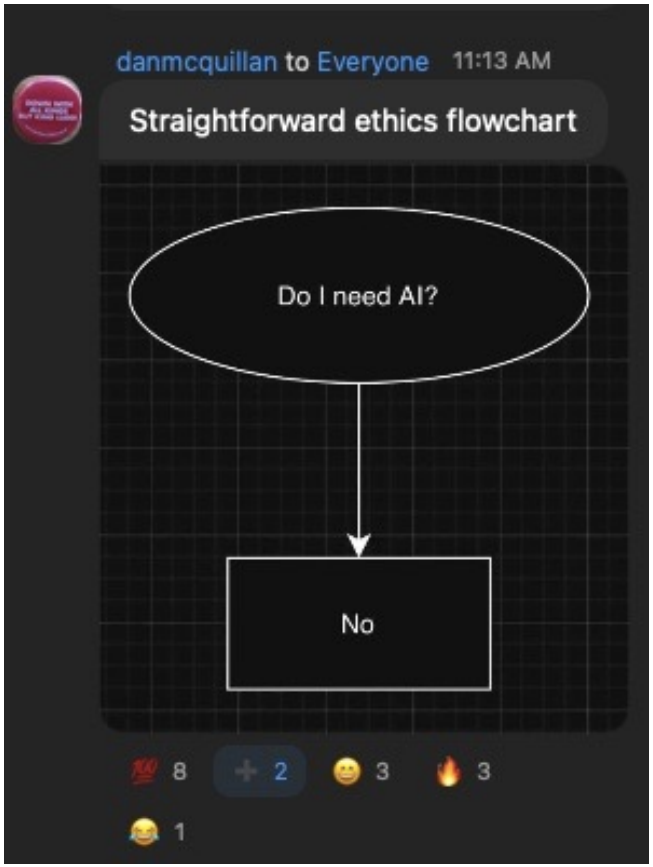

Dan McQuillan’s AI ethics flowchart, from LIMITS chat, June 24

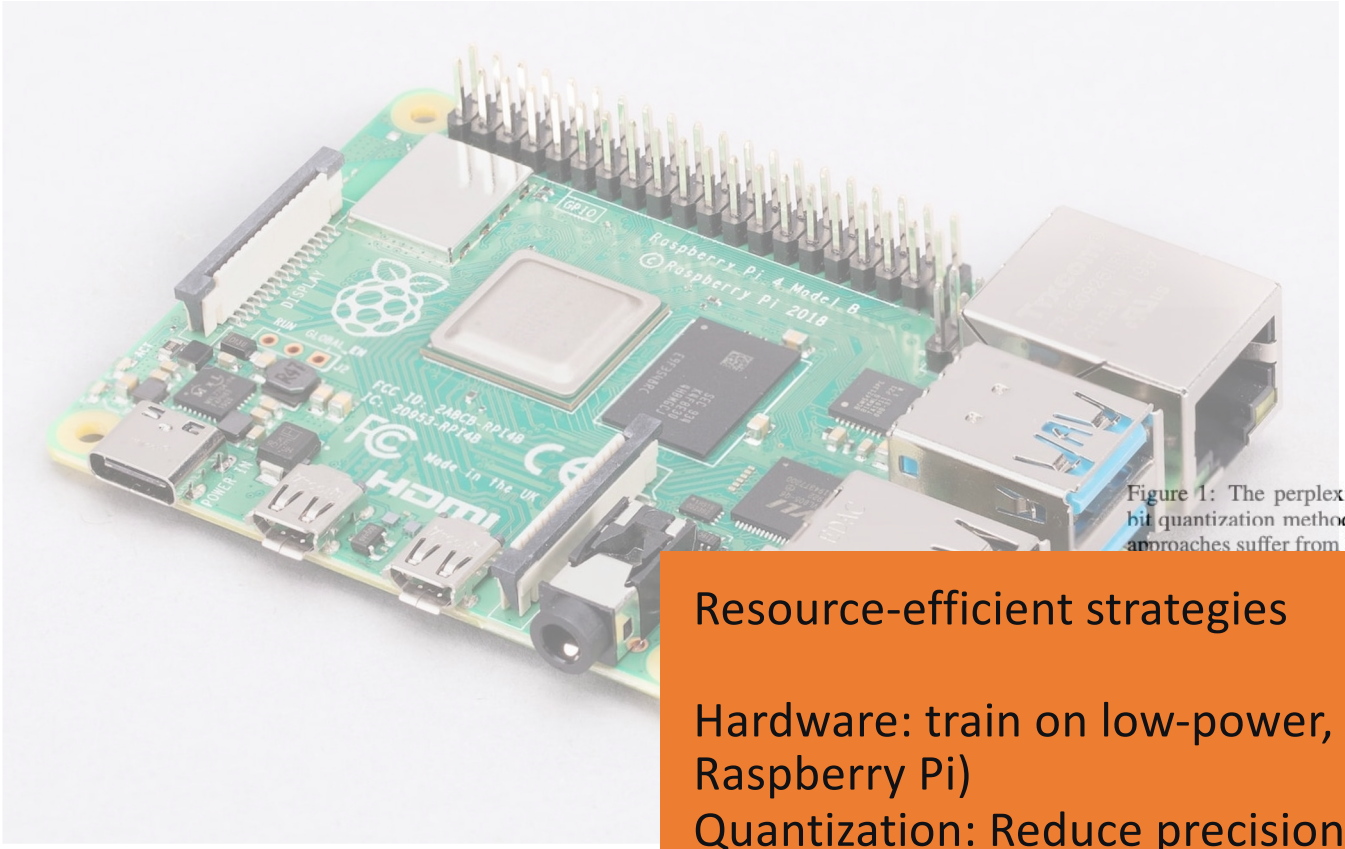

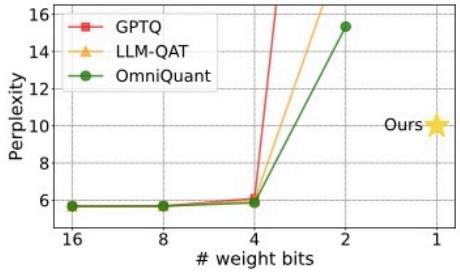

Figure 1: The perplexity (lower scores mean better performance) of existing widely-used low-bit quantization methods on LLaMA-7B, reported on Wikitext2 [23]. All the examined previous approaches suffer from significant performance degradation when quantizing models to 2-bit values.

Resource-efficient strategies

Hardware: train on low-power, consumer-grade hardware (e.g. Raspberry Pi)
Quantization: Reduce precision of numerical values (using fewer bits) to shrink model size while maintaining overall accuracy
Data-efficient training: Leverage few-shot learning techniques to enable generalization from a smaller dataset (Xu et al. 2024)

Generating novelty

- Noise: Add noise to image embeddings rather than working with averages, to introduce surprise and randomness and avoid generating "mean images” (Jess)
- Activation functions & parameters: Different activations produce different styles; parameters control specific style aspects. Slight changes lead to slight changes in the image, while larger changes lead to more significant and unpredictable alterations to both structure and style.
- Base Model: Adapt StableDiffusion v2-1-unclip (small) — adding noise and surprise to image-plus-word prompts, heavily weighted with our own ethically sourced image database.
- Dataset Weighting (Jess)

Artistic approaches to tiny image-generation models

- Novelty. The goal of this tiny image generator is novelty, not averages, and to magnify artifacts
- The model will reflect the community that uses it
- Victoria & Albert Museum dataset of textiles and Islamic calligraphy
- What other ethically sourced datasets could we use?
- Dataset weighting: different data sources are assigned different levels of importance during model training or fine-tuning. Counters biases to photorealism and Western art
- Unweighted training gives generic results

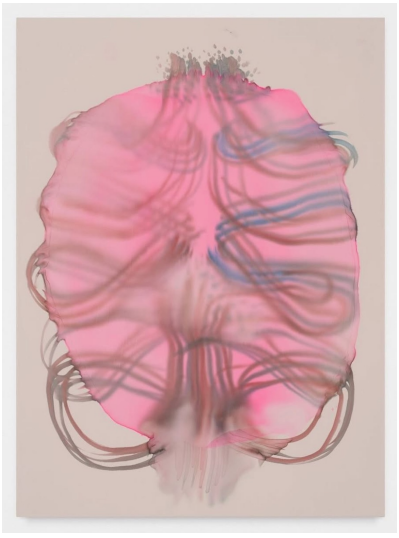

LLM have statistical biases built in creating what Hito Steyerl calls “mean images”. Mean here means both average and exclusionary.

FIGURE 2: *Stable Diffusion’s ‘Image of Hito Steyerl’*

Source: stablediffusionweb.com

Artistic approaches to tiny image-generation models

- Novelty. The goal of this tiny image generator is novelty, not averages, and to magnify artifacts
- The model will reflect the community that uses it
- Victoria & Albert Museum dataset of textiles and Islamic calligraphy
- What other ethically sourced datasets could we use?
- Dataset weighting: different data sources are assigned different levels of importance during model training or fine-tuning. Counters biases to photorealism and Western art
- Unweighted training gives generic results

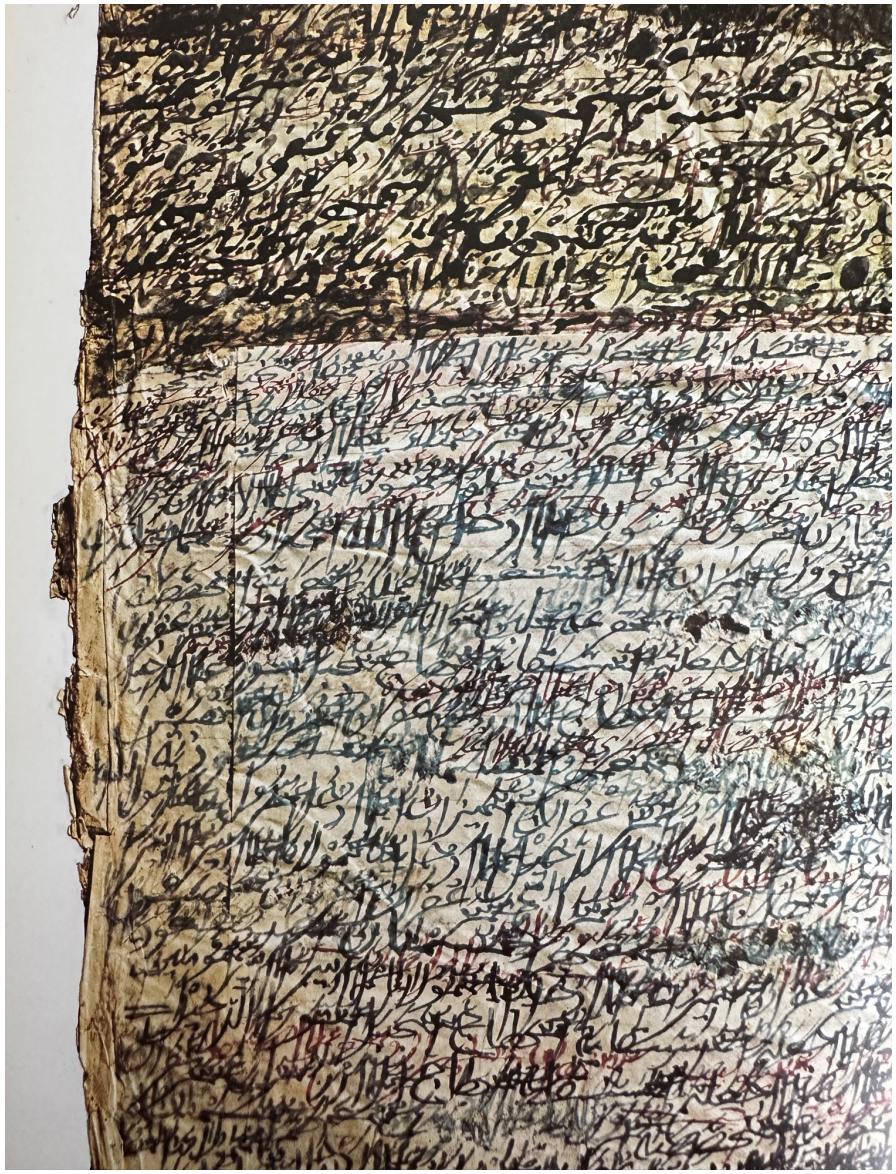

A political element

- Most AI has been developed in the global North and therefore reflects hegemonic interests.
- Ongoing dehumanizing and Islamophobic portrayal of Palestinians allows for the continuation of genocide.
- Israel makes use of and is invested in AI-based facial and emotion recognition technologies, automated weapons, social media monitoring, and military targeting, and it has used machine learning to generate targets for bombing in Gaza

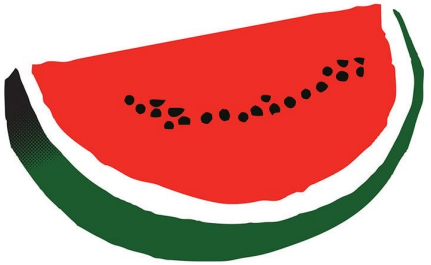

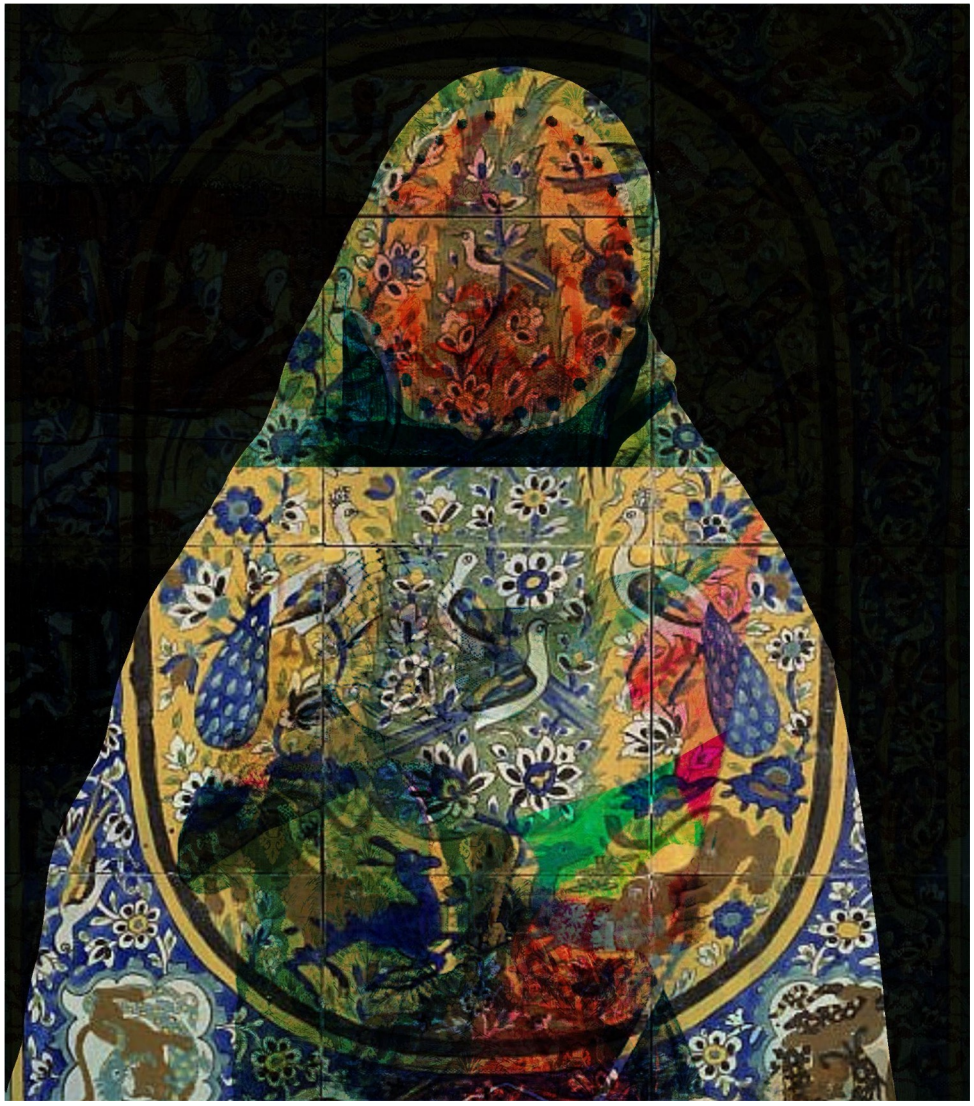

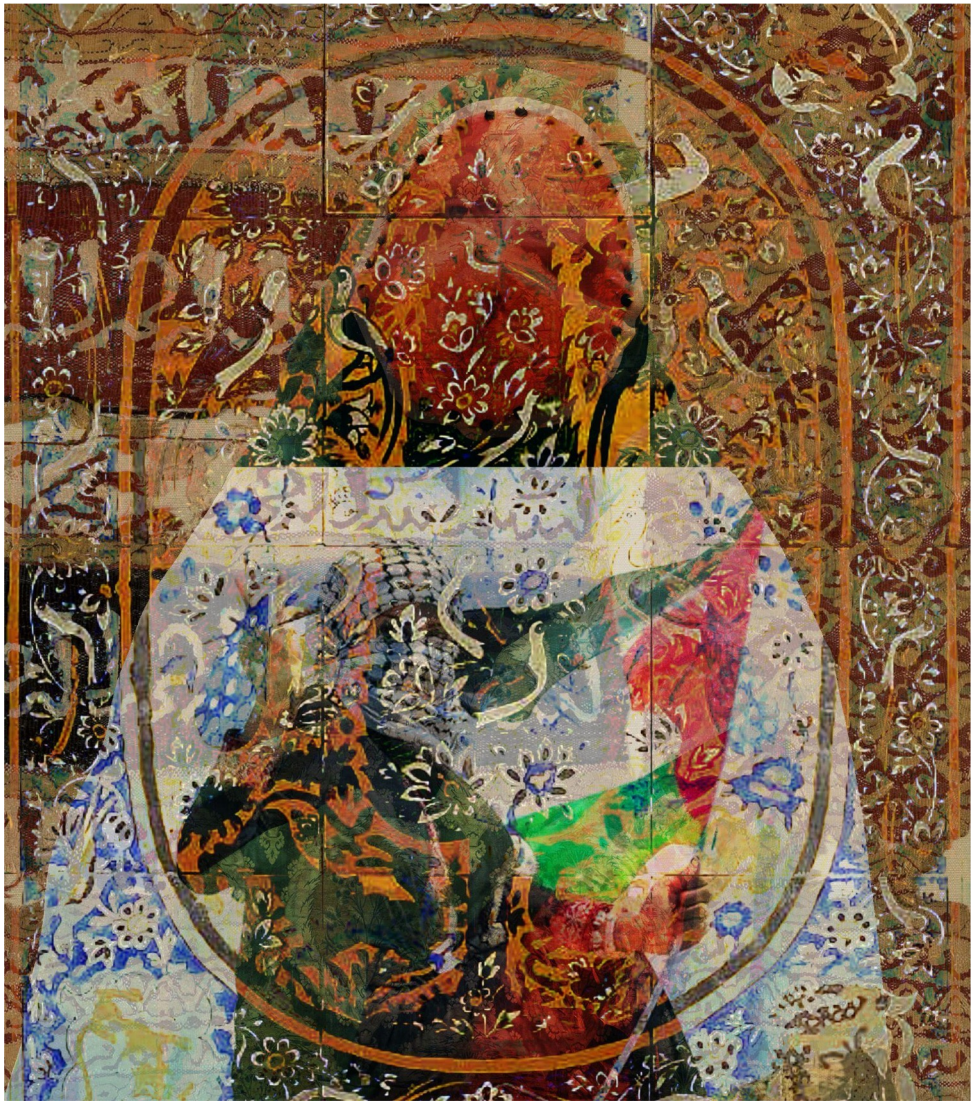

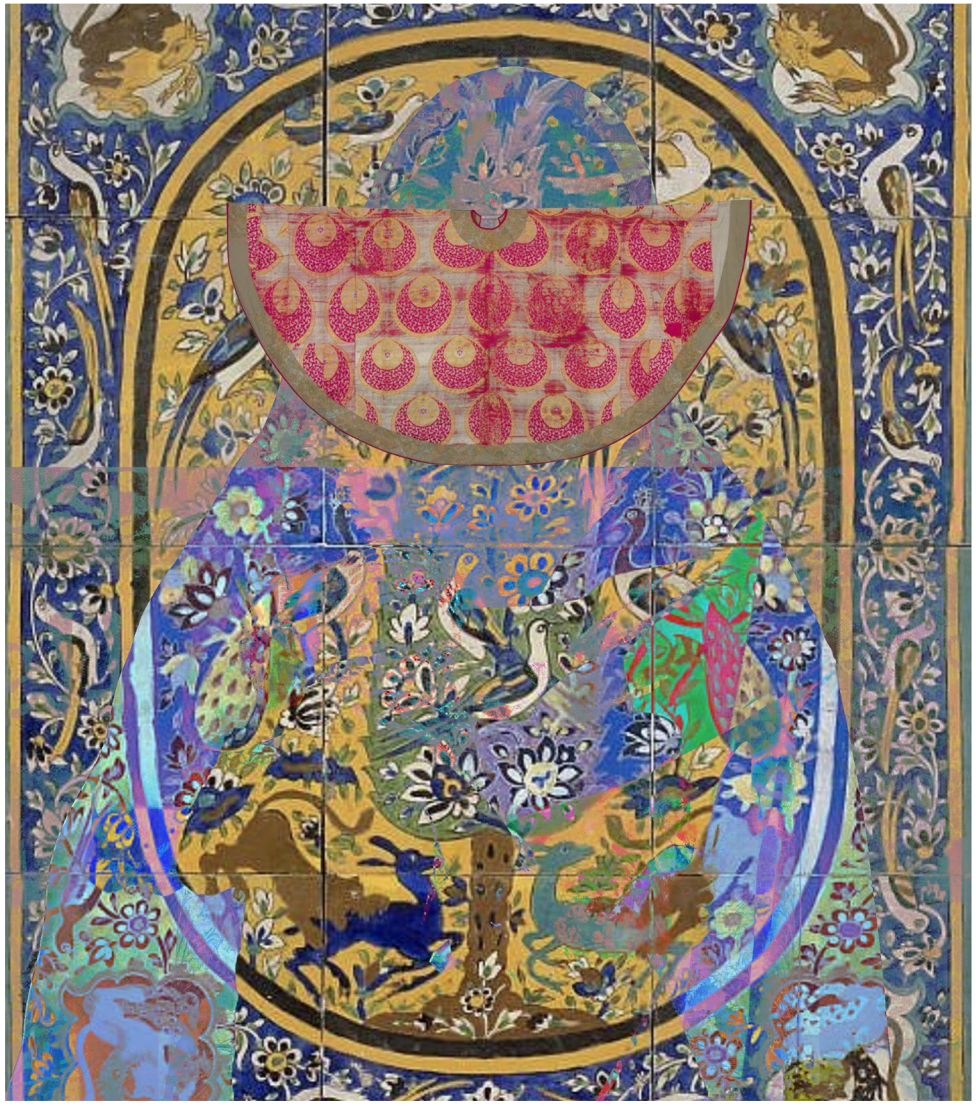

Thank you!

lmarks@sfu.ca
jam47@sfu.ca
kehui_li@sfu.ca