From Pixels to Metal: AI-Empowered Numismatic Art

Penousal Machado[1] , **Tiago Martins**[1] , **João Correia**[1] , **Luís Espírito Santo**[1,2] , **Nuno Lourenço**[1] , **João Cunha**[1] , **Sérgio Rebelo**[1] , **Pedro Martins**[1] , **João Bicker**[1]
[1]University of Coimbra, CISUC/LASI, DEI, Coimbra, Portugal
[2]Vrije Universiteit Brussel, Brussels, Belgium
{machado, tiagofm, jncor, lesanto, naml, jmacunha, srebelo, pjmm, bicker}@dei.uc.pt

Abstract

This paper describes our response to a unique challenge presented by the Portuguese National Press-Mint: to use Artificial Intelligence to design a commemorative coin that celebrates the "digital world." We explain the process of this coin's co-creation, from conceptualisation to production, highlighting the design process, the underlying rationale, key obstacles encountered, and the technical innovations and developments made to meet the challenge. These include the development of an evolutionary art system guided by Contrastive Language–Image Pre-training (CLIP) and Machine Learning-based aesthetic models, a system for prompt evolution, and a representation for encoding genotypes in mintable format. This collaboration produced a limited edition 10 euro silver proof coin, with a total of 4 000 units minted by the National Press-Mint. The coin was met with enthusiasm, selling out within two months. This work contributes to Computational Creativity, particularly co-creativity, co-design, and digital art, and represents a significant step in using Artificial Intelligence for Numismatics.

1 Introduction

In recent years, the application of Artificial Intelligence (AI) for creative and artistic endeavours has attracted considerable attention, increasing the opportunities to use or misuse AI for various art and design tasks. One such opportunity arose when INCM, the Portuguese National Press-Mint, invited us to use AI to design a commemorative coin themed around the "digital world."

Coins are legal tender and also enduring symbols of cultural heritage. They are distinct from commemorative medals by their dual role as aesthetic artefacts and practical currency. While medals commemorate events, coins hold designated monetary values and allow economic transactions under governmental authorisation. Anyone can mint a medal, but only the government can mint a coin.

The duality of coins as objects of utility and beauty, along with practical considerations such as *mintability*, durability, recognisability, legal requirements, thematic significance,

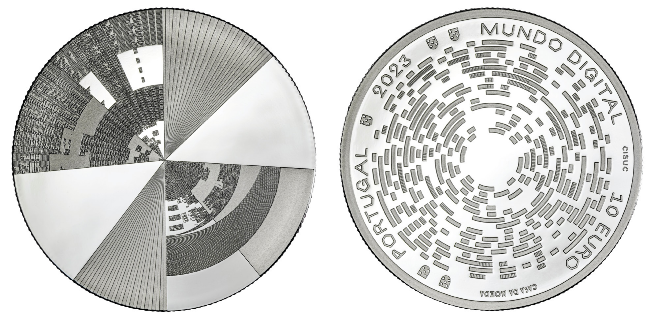

Figure 1: Reverse (left) and obverse (right) of the coin "Digital World." Photos by Nuno Silva / INCM.

and aesthetic appeal, present unique challenges for their design. Taking these requirements into account, we aimed to produce a coin that honours the rich tradition of numismatics and of the National Press-Mint, while respecting our own identity as an AI and Design Lab, and respecting our views about Computational Creativity (CC) research, aiming to challenge traditional boundaries and offer novel avenues for expression and design.

As such, the goal of this project was not only to create a coin themed around the "digital world" but also to explore how co-creativity can contribute to expanding traditional numismatic artistry, bridging the gap between the coin minting heritage and the essence of computational aesthetics. It also highlights AI's potential to act as a co-creator in complex design processes, showcasing how CC can avoid pastiches and stereotypical images, often associated with modern generative AI.

The novelty of our work lies in (i) the novel AI tools developed during this project, namely a new Evolutionary Computation (EC) engine that takes advantage of the semantic interpretation capabilities of CLIP [Radford *et al.*, 2021], Machine Learning (ML) aesthetic models and an EC system for evolutionary prompt design; and (ii) the application of these to the historically and culturally rich field of numismatics.

The culmination of this project was the minting of a limited edition (4 000 exemplars) 10 euro silver proof coin by INCM (Figure 1), which sold out in less than two months at a fixed price of 95 euros each, attesting its success and the public's approval of the design and inherent design principles. Through this work, we aim to contribute to the discourse on the role of AI in creative industries, offering insights and methodologies that could inspire future projects at the intersection of technology, art, and tradition.

The remainder of the paper is structured as follows: in Section 2, we discuss the literature that is related to the approach, taking into account its multidisciplinary aspects. In Section 3, we contextualise the project, describing its origins and presenting preliminary experiments. In Section 4, we detail the methodology used to generate the reverse face of the coin, and in Section 5, we explain the method used to encode the genetic code of the reverse on the obverse face. Finally, Section 6 presents unique digital artefacts generated by our system, and Section 7 gathers the main conclusions and points towards future work.

2 State of the Art

A project of this nature is imminently multidisciplinary, involving several scientific domains, including CC, Co-Creativity, Design, AI, ML, EC, and Generative AI. More than providing an in-depth review, we wish to contextualise our project within the broader discourse of CC and the role of AI in art and design.

CC can be defined as the area of understanding and developing AI systems that can perform tasks that would be considered creative if carried out by humans [Wiggins, 2006]. Boden's seminal work [2004; 2009] provides a basis for understanding AI's role in generating novel and valuable ideas or artefacts, establishing the foundations for CC research. Current models of the creative act, such as the formalisation proposed by Wiggins [2006], later refined by Colton and Wiggins [2012], and the FACE and IDEA descriptive models proposed by Colton et al. [2011] further refine our understanding of the creative process. Aspects such as novelty, value and intentionality have been considered as indicators of creative behaviour [Ventura, 2016], with authors proposing different methodologies to assess creativity in computational systems (see, e.g., Jourdanous [2012]). Recently, notions from CC have been employed to assess the creativity of deep generative approaches [Berns and Colton, 2020; Oppenlaender, 2022], questioning if they can be considered creative to some degree [Brown and Jordanous, 2022] or "mere generators" [Ventura, 2016].

Although Generative AI has a long history, as illustrated by the works of Cohen [1995], Cope [2005] or Sims [1991], to name a few, the development of Generative Adversarial Networks (GANs) [Goodfellow *et al.*, 2014] was a turning point in the use of AI for image generation. To the best of our knowledge, Elgammal et al. [2017] "Creative Adversarial Networks" are the first to apply GANs for CC purposes by exploring the idea of using GANs to produce art that deviates from established norms and styles. It is worth noting that before the advent of GANs, Machado et al. [2008] had already presented a system with a GAN-like architecture for artistic purposes, although with an evolutionary system as the generator.

The development and refinement of ML models, including Variational Autoencoders (VAEs) and, more recently, Transformers and Diffusion Models, have further expanded the horizons of Generative AI [Karras *et al.*, 2019; Vaswani *et al.*, 2017; Ho *et al.*, 2020], giving rise to popular projects such as DALL-E [Ramesh *et al.*, 2021], Stable Diffusion [Rombach *et al.*, 2022], and Midjourney [Hachman, 2022] that demonstrate the ability of data-driven AI approaches to produce complex, context-rich images from textual descriptions.

The field of human-AI collaboration, especially in art and design, focuses on demonstrating that AI systems can take the role of co-creator, assisting and promoting human creativity. Co-creation is arguably augmenting human creativity through the generation of novel ideas [Du Sautoy, 2019; Eiben and Smith, 2015] but also led to the development of computer-assisted creativity systems where humans and AI systems jointly produce artworks and designs, possibly redefining the conventional roles of artists and designers. Works about co-creative processes emphasise the importance of understanding human-agent interaction dynamics [Jordanous, 2017; Espírito Santo *et al.*, 2023], while others explore various dimensions of co-creativity, from participatory sense-making and mixed-initiative co-creativity to the prospect of autonomous co-creative agents [Davis *et al.*, 2016; Yannakakis *et al.*, 2014].

EC, particularly in the form of evolutionary art, has been used to automate creative expression. By employing algorithms that simulate the process of natural selection, these systems generate artworks that evolve based on a set of aesthetic criteria. In the design domain, EC has been successfully applied to produce several kinds of artefacts, such as typefaces [Martins *et al.*, 2018; Parente *et al.*, 2023], imagery [Machado *et al.*, 2014; Cunha *et al.*, 2019], posters [Rebelo *et al.*, 2024], etc. For further examples see [Romero and Machado, 2008; Veale and Cardoso, 2019; Machado *et al.*, 2021]. Recent advancements include automated fitness assignment methods and the integration of ML techniques into the evolutionary loop, enhancing the diversity and complexity of the outputs (see [Banzhaf and Machado, 2024; Correia *et al.*, 2024] for recent surveys on the topic).

In the field of numismatics, AI is primarily used for classifying coins (e.g., see Anwar et al. [2021]). The Pressburg Mint [2023] introduced a coin designed by AI. No detailed study on the design process is available, yet by inspection of the final product, we suspect it results from an off-the-shelf model such as DALL-E or Stable Diffusion. It is unclear whether this coin was designed before or after our project; however, it was minted before. Regardless, we were unaware of it until recently, and its design approach does not align with our perspective on computational creativity and art.

3 Design Brief and Implications

This project started in May 2022, following a unique enquiry from INCM: "Can you create a coin themed around the digital world using Artificial Intelligence?"

This question sparked our collaboration with INCM and prompted us to explore essential concepts: the essence of a coin, the distinction between a coin and a medal, the intended recipient of an AI-designed coin, and the nature of a coin designed by AI.

Figure 2: Examples of coins themed around the digital world designed by AI generated with off-the-shelf generative AI models.

Focusing on these last two questions, it became obvious that if the AI was designing the coin for itself, it would not be a physical coin at all; most likely, it would be something incomprehensible to humans. Although this concept is exciting, the boundaries of the design challenge imply that the coin must be mintable and appealing to human consumers. Therefore, the AI needed to design a coin for humans. This realisation led to our primary design principle:

1. The coin should represent AI's interpretation of a coin related to the "digital world."

Subsequently, we outlined further specifications: it would be a "classical" 10 euros commemorative coin, crafted in silver, with a traditional round shape. We decided against using colour or unconventional shapes to maintain authenticity and avoid gimmicks.

During our meetings with INCM, we realised that although our lab included several renowned designers, designing coins required specific knowledge. In one of these meetings, the INCM director explained, "A coin has two faces, and in a good coin, these faces should communicate with each other."

This resonated with our views regarding AI and CC's goals and purposes, namely with our posture towards human-centred AI and human-machine collaboration, leading to the definition of two other design guidelines:

2. One side of the coin is designed by AI for humans and the other by humans for AI.
3. Both sides should convey the same concept, using a representation tailored to the intended audience of each side and preferably appealing to both.

Ultimately, four thousand silver-proof coins with a 925‰ fineness were minted, each with a 40mm diameter and a 10 euro face value. The coin was officially released on December 11, 2023, at the Imprensa Nacional–Casa da Moeda in Lisbon. It was available at INCM's physical stores or online shop and eventually sold out by the end of January.

3.1 Early Experiments

Given the popularity of deep learning generative models such as Stable Diffusion, Midjourney and DALL-E, our initial experiments involved using these tools to create "A coin themed around the digital world designed by AI." These attempts yielded predictable and disappointing results (see Figure 2). The coins featured clichéd and derivative imagery, often depicting AI as the profile of a female face with a digital circuit on the brain, typically facing left. We decided this was not what we envisioned, nor did it align with the perspective on Computational Creativity and AI that we support.

3.2 Refining the Design Brief

We realised that the coin could not be designed by just any AI; it must be an AI created by us and aligned with our views. Likewise, it is not just any coin; it is a coin that will belong to the collection of INCM. Thus, we outlined two new design guidelines as follows:

4. The AI must be aligned with our views. It is our lab's next project. It must fit.
5. The AI must be aligned with the INCM collection. It is the next coin of the collection. It must fit.

These refined guidelines allowed us to proceed with the development in a way that is aligned with our views. The subsequent sections detail the development of both sides of the coin, describing the technical innovations.

4 By AI for Humans: Generating Coins Themed Around the Digital World

This section details the development and application of an evolutionary art system to create the coin's design. First, we describe the process of evolving coin heightmaps using Genetic Programming (GP). Then, we discuss the integration of tools like CLIP and ML aesthetic models to guide the evolutionary process. Finally, we present the results of our experiments and the design choices made to achieve a visually captivating and thematically appropriate coin.

4.1 Evolving Coin Heightmaps

Following the expression-based evolutionary art approach popularised by Sims [1991], we use GP to evolve a population of images. Each individual's genotype is a mathematical expression, and the genotype-to-phenotype mapping involves evaluating these expressions to generate images. This is done by calculating the output of the expression for every pixel in the image space, with the expression's variables x and y corresponding to the pixel's coordinates, and the expression's output determining the pixel's colour.

We use pyNEvAr[1], a modern version of our evolutionary art system NEvAr [Machado and Cardoso, 2002], implemented using TensorGP [Baeta *et al.*, 2021], which enhances traditional tree-based GP by incorporating tensor operations with the TensorFlow Python library, using GPUs for faster evaluation of individuals. Additionally, TensorGP improves efficiency by caching intermediate results, reducing the redundancy of computing common sub-trees.

Most notably, our evolutionary engine can communicate with CLIP [Radford *et al.*, 2021], and in addition to its own aesthetic model, it has access to publicly available ones. Specifically, we can specify a text prompt, such as "a coin themed around the digital world," and use CLIP to convert this prompt into a text feature vector. For each image generated by the evolutionary system, we use CLIP to extract the associated image feature vector. By calculating the cosine similarity between the text feature vector and each image feature vector, we can assign fitness to the images, allowing us to evolve images that match the given prompt. Additionally,

[1]For further information visit: github.com/cdvetal

Parameter	Value
Generations	100, 250, 500
Population size	100
Elitism (elite size)	1
Tournament size	5
Mutation probability	0.9
Mutation operators	delete, insert, point, subtree
Crossover probability	0.5
Crossover operator	random sub-tree swap
Minimum initial depth	1
Maximum initial depth	6
Minimum allowed depth	1
Maximum allowed depth	12
Bloat control	weak depth control [Silva and Costa, 2009]
Generation method	Ramped Half and Half
Fitness metric	Cosine similarity
Function set	abs, cos, exp, frac, log, neg sin, sqrt, tan, add, and, div, max, mdist, min, mult, or, pow, sub, xor, ifthenelse, warp
Terminal set	α, r

Table 1: Settings of pyNEvAr for the presented experiments.

since the system can interact with ML aesthetic models like LAION-5B [Schuhmann *et al.*, 2022], it can use these outputs to evolve images that maximise aesthetic appeal according to the chosen model.

Given that we established that the coin would be circular, polar coordinates were used instead of Cartesian ones; as such, the variables x and y were replaced by angle, α, and radius, r. The images produced are interpreted as heightmaps, creating the 3D surface of one of the faces of the coin. Figure 3 exemplifies the genotype-to-phenotype mapping process, presenting three expression and image pairs. Table 1 presents the general settings of the evolutionary algorithm, including the function and terminal set.

One remarkable difference between our system and popular generative AI tools is that the generator is not data-driven; it has never seen a painting by Van Gogh, Picasso, or any other painter. It also does not use or imitate human tools and design primitives. There is no point, line or plane. It is designed with innovation in mind, not imitation.

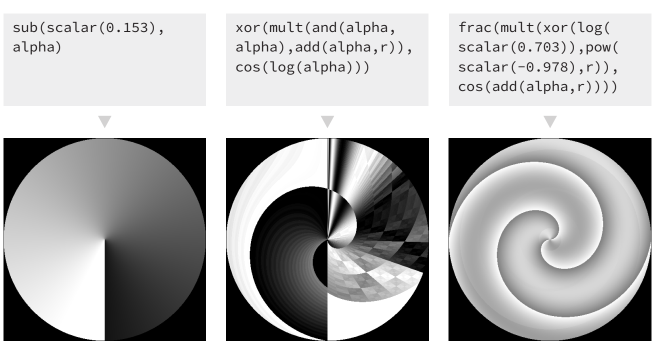

Figure 3: Examples of genotype-phenotype pairs.

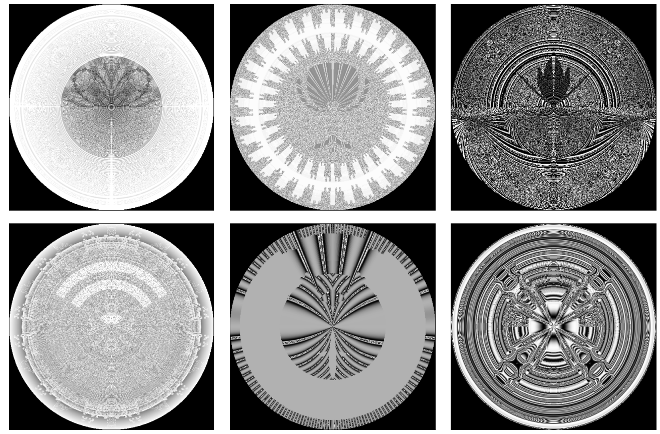

Figure 4: Examples of images evolved while trying to match the prompt “a coin themed around the digital world.”

To some extent, a prompt such as “a coin themed around the digital world” works, but the images tend to be mediocre (see Figure 4 for examples). A perusal of the result reveals several issues that prevent high-quality results; for instance, the word “world” becomes an “attractor” that leads to the creation of a circular pattern in the middle of the image; the images tend to be low contrast because most photos of coins are also low in contrast; in some cases, the images appear cropped.

We could do prompt engineering, but we would be driving evolution, and we want the coin to be designed by AI, not by us. To avoid such bias, we developed a new tool that would allow us to evolve prompts, which we summarily describe in a later subsection.

According to our design guidelines, the system should produce a coin that fits INCM’s collection and aligns with our group’s work (see Subsection 3.2). Therefore, it needs to be aware of the past production of both INCM and our group. Since our generator is not data-driven, using the prompt becomes the best way to ensure these contexts are considered when assigning fitness.

Incorporating the Context

To incorporate context, we created an image database composed of contemporary INCM coins, analysed 910 INCM coins, and, with the help of their experts, selected 34 coins according to practical criteria such as the existence of high-quality frontal photos without reflections but also the content of the coin, e.g., a coin portraying the face of a famous person is not particularly valuable for our approach. Using CLIP interrogator [pharmapsychotic, 2023], we extracted roughly 400 expressions from these coins and curated them to remove those related to content (e.g., flower, ship, horse, guitar, building, men, king).

We applied the same process using our pyNEvAr. We curated high-quality images evolved without human intervention but also images developed for this task using a mix of user-guided evolution and automated evolution. In the automated approach, fitness was assigned to maximise the aesthetic score of LAION-5B [Schuhmann *et al.*, 2022]. Fig-

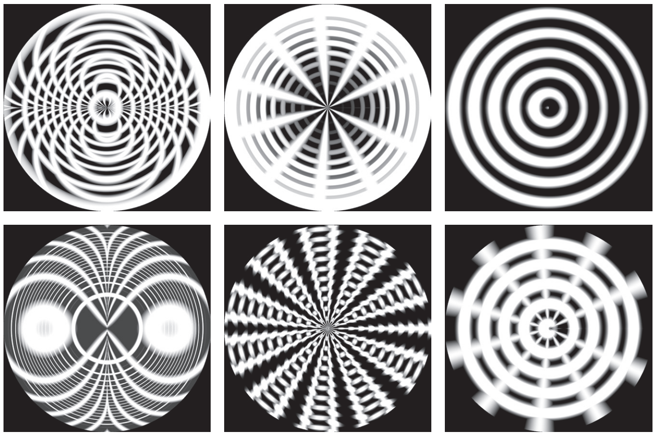

Figure 5: Examples of images evolved while attempting to maximise the aesthetic score of LAION-5B.

ure 5 shows some examples of these images. As previously, we used CLIP interrogator to extract expressions from these images, containing a series of unexpected yet understandable expressions that appear to enhance image quality. Some notable examples include: "top view", which avoids non-perpendicular images of coins; "uncompressed png", which promotes images with clean transitions from black to white; "circular object", which avoids cropped images. We eliminated all expressions related to specific styles (e.g., "in the style of Escher") and content (e.g., "eyes").

EC Engine for Prompt Evolution

The system for prompt evolution developed specifically for this project is called MetaPrompter [Martins *et al.*, 2023]. It allows users to: (i) design a *template for a prompt* that defines the search space of various alternative prompts, and (ii) interactively evolve a population of prompts that explore different points within this space according to their preferences. This provides a dynamic and interactive method for finding and fine-tuning prompts, enabling users to obtain images that align with their preferences.

Designing image generation prompts is streamlined through a structured method that organises elements like subject, verb, and object, allowing for dynamic and static components. Dynamic components are specified within angle brackets and can include multiple terms separated by vertical bars, creating various prompt variations from a single template. This system enumerates possible terms for each component, which can be directly included in the meta prompt or linked to external lists for greater flexibility. It also allows combining terms within elements, with options to specify the combination count and repetition. Furthermore, components can be repeated within prompts to emphasise certain aspects. This approach enhances creative possibilities for generating images by enabling the dynamic creation and modification of prompt components and their combinations. More information can be found in Martins et al. [2023].

Figure 6 presents the prompt template we used, containing a mix of fixed and dynamic components. Based on the previous analysis, we decided that the terms "black and white

```
meta_prompt ::= "black and white image, an image of a <DIGITAL>
                <WORLD:0-1>, a <RAYTRACED> image, inspired by
                <ARTIST>, <in the style of STYLE:1-2:>,
                <ABSTRACT>, top-view"

<DIGITAL>   ::= get_related_terms("digital",
                include_input_term=True, max_results=10)

<WORLD>     ::= get_related_terms("world",
                include_input_term=True, max_results=10)

<ABSTRACT>  ::= get_related_terms("abstract",
                include_input_term=True, max_results=10)

<ARTIST>    ::= "Pedro Pedraja" | "Gilberto Soren Zaragoza" |
                "Adreas Rocha" | "Luis Royo" | ...

<STYLE>     ::= "3d" | "aboriginal art" | "abstract artwork" |
                "acrylic artwork" | "acrylic marbling art"| ...
```

Figure 6: MetaPrompter grammar

image" and "top-view" should always be present. Then, we added an expression to define the image's content: "an image of a <DIGITAL> <WORLD:0-1>", with two dynamic terms, DIGITAL and WORLD. When an evolutionary run is initiated, the system consults the "Related Words" API[2] that combines several different algorithms which compete to get their results higher in the list (e.g., ConceptNet and WordNet), retrieving a fixed number of concepts in the vicinity of the one provided. Thus, the term WORLD can be replaced by "world", "planet", "earth", etc. The next two expressions "inspired by <ARTIST>" and "<in the style of STYLE:1-2:>" provide stylistic clues. The term ARTIST can be replaced by one of the artists of the corresponding list. Notice that this list results from the expressions extracted by CLIP interrogator from the existing INCM coins. Likewise, the term STYLE can be replaced by an element of the STYLE list, which was also extracted by CLIP interrogator from the images evolved by our lab and INCM coins. Finally, we include the term ABSTRACT to promote abstract imagery; as before, this term can be replaced by related concepts.

Using the MetaPrompter interface, we evolved several satisfying prompts, able to produce high-quality results and selected 14 to conduct the evolutionary runs of coins, with three examples being: (i) "black and white image, an image of a broadcasting world, inspired by Alvar Aalto, in the style of art nouveau, in the style of figurativism, abstract, top-view"; (ii) "black and white image, an image of a internet, inspired by Ai Weiwei, society, in the style of art nouveau, behance contest winner, simulacra and simulation, top-view"; and (iii) "black and white image, an image of a broadcasting, inspired by Nadir Afonso, in the style of massurrealism, in the style of figurativism, futurism, ethereal, top-view."

4.2 Experimental Results

We conducted several evolutionary runs using the evolved prompts to guide the process. We performed 30 independent runs for each of the 14 selected prompts, totalling 420 runs and evolving roughly 43 million images. A process of automatic curatorship, inspired by the work of Correia et al. [2019], reduced this number to 1 974 images. The process can be summarised as follows: We start with an empty global

[2]https://relatedwords.org/

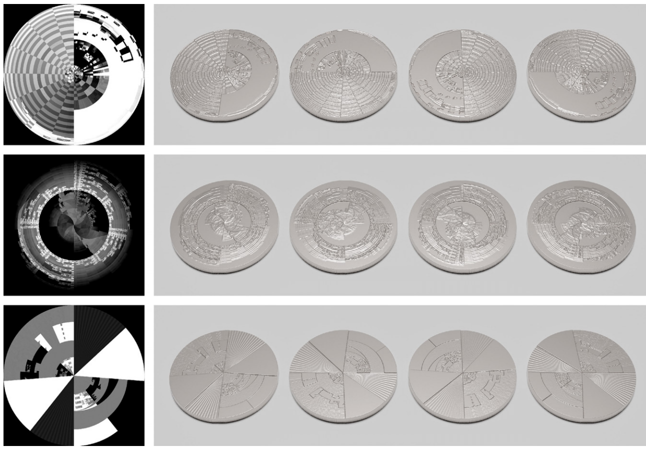

Figure 7: The 3 evolved images (left) and corresponding 3D renderings (right) selected by the creative director of our lab. The bottom row shows the coin design unanimously selected by the team.

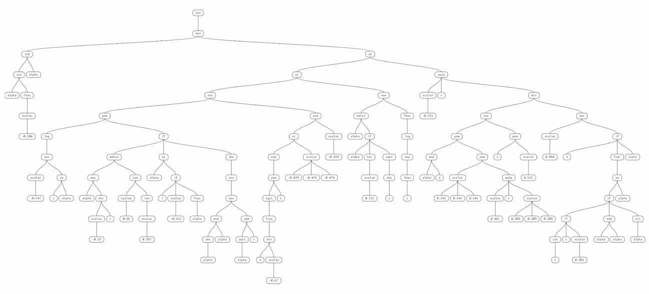

Figure 8: Visualisation of the expression that generates the coin's reverse face. The expression tree rendering is available at: https://cdv.dei.uc.pt/2024/coin-expression-visualisation.mp4

archive of images. For each run, we add the fittest image to the archive. The fittest images from each population are compared with those already in the archive and are added if they are sufficiently different from the stored images.

The 1 974 images were then reduced to 142 through visual inspection. These images were 3D rendered, and a blind vote by our team led to a selection of 24 images. The creative director of our lab selected three, and the team unanimously picked one. Figure 7 shows the three images (and the corresponding 3D renderings) selected by the creative director, where the third one (bottom) was the one selected by the team. Figure 8 presents the genotype of the selected individual and a link to a video that illustrates how the different subtrees contribute to the final result.

4.3 The Coin

After selecting the image, we inspected the prompt that was used to generate it: "black and white image, an image of a broadcasting world, inspired by Alvar Aalto, in the style of art nouveau, in the style of figurativism, abstract, top-view."

Its analysis reveals that "Alvar Aalto" is a keyword extracted from INCM coins. Interestingly, this keyword comes from the coin designed by Eduardo Souto de Moura to honour Siza Vieira's work, both of whom are Pritzker Prize-winning

Figure 9: The coin designed by AI themed around the "Digital World" immediately after minting (reverse face). Photo by Nuno Silva / INCM.

architects. Alvar Aalto significantly influenced Siza Vieira's work, and although the CLIP interrogator did not identify the names Eduardo Souto de Moura and Siza Vieira, the inclusion of Alvar Aalto's name contributed to the architectural nature of the result.

A human not aware of the process can interpret it as follows: "The image is divided into 6 circular sectors and 3 layers of distinct visual style: a first smooth layer, the surface; a second with abstract structures and emerging patterns; a third layer of concentric lines with two depth levels, which form a binary base" — *INCM*, 2023.

The last step was getting the INCM's approval, not only regarding the design's aesthetic and conceptual integrity but also regarding its practicality for reproduction and minting. Despite initial reservations about the design's feasibility, the skill and expertise of the INCM minting team ultimately made it possible. Figure 9 shows the coin immediately after minting, while Figure 10 displays both sides of the coin in an angled shot, highlighting the level of detail.

Since we decided that the facial value of the coin and other required details would be on the side designed by humans, this face became the reverse side. The design of the obverse is described in the next section.

5 By Humans for AI: Code Representation

The coin's obverse face is a graphical representation of the genetic code of the reverse image (see Figure 1). To achieve this, we developed a system that converts the genotypes evolved by AI into a visual composition of concentric arcs, encoding the mathematical expression corresponding to the reverse image. This system uses a dictionary-based compression algorithm to encode the sequence of functions and values that constitute the genotype. Continuous arcs represent mathematical functions, while interrupted arcs represent numerical values. Inspired by data storage disks, this arc-based system is designed to be machine-readable, providing access to the instructions needed to reproduce the reverse image. Surrounding the genotype representation, we added graphic

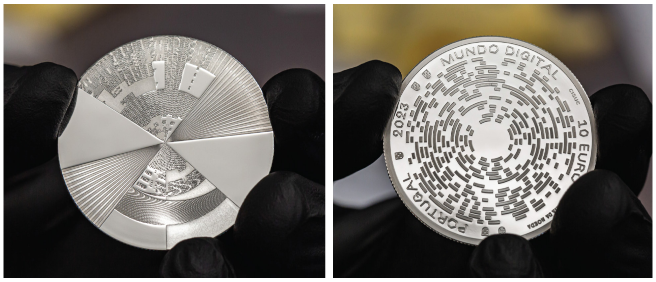

Figure 10: Reverse (left) and obverse (right) faces of the coin "Digital World." Photos by Nuno Silva / INCM.

elements and typography required for official currency. The typography was specifically generated for the coin using the parametric system LetterSpecies [Pereira *et al.*, 2019].

6 Unique Digital Assets

Although buyers were unaware, we created a unique digital artefact for each of the 4 000 coins, which includes a "digital offspring" of the coin. These digital assets illustrate the evolutionary process that led to the minted coin. Each digital asset is a 40-second video showing the evolutionary process behind the coin's creation, divided into three segments.

The first part of the video represents the ancestors of the image used on the coin's reverse side. For this purpose, we provided a brief overview of the evolutionary process, by creating a genetic morphing [Sims, 1991] animation performed between a series of individuals selected from the ancestors of the image. In the second part, the image transforms into a 3D coin model, symbolising the minting process. The 3D coin begins to rotate, showcasing its two faces. The transition to the third and final part of the video occurs subtly during the coin's rotation. This part presents a digital offspring, i.e., a new coin created by continuing the evolutionary process for a few extra generations. In this case, the evolutionary process is entirely guided by LAION-5B [Schuhmann *et al.*, 2022]. We use a process similar to the one described earlier in Subsection 4.2 to ensure that each asset is unique. We generated 14 000 offspring of the original coin, sorted them by aesthetic rating, and added them to an initially empty archive starting from the fittest. We ensured diversity by comparing genotypes and phenotypes with existing archive images. If similar images already existed, only the fittest one remained in the archive. This allowed us to create a different digital asset for each of the 4 000 coins. Figure 11 shows six of these assets rendered with different materials.

Rarity was introduced by varying the materials used in rendering the descendant coin displayed in the final segment. There are five rarity levels, with 2 400, 1 000, 400, 160, and 40 digital assets. In ascending order of rarity, the descendant coin, unique to each asset, is rendered in nickel, copper, rose gold, gold, and chameleon (see Figure 11 and video at https://cdv.dei.uc.pt/2024/digital-asset-0001.mp4). Additionally, there are 40 extra digital assets to be awarded to collectors who possess at least one of each rarity level. To create these, we used the same process but allowed the use of colour.

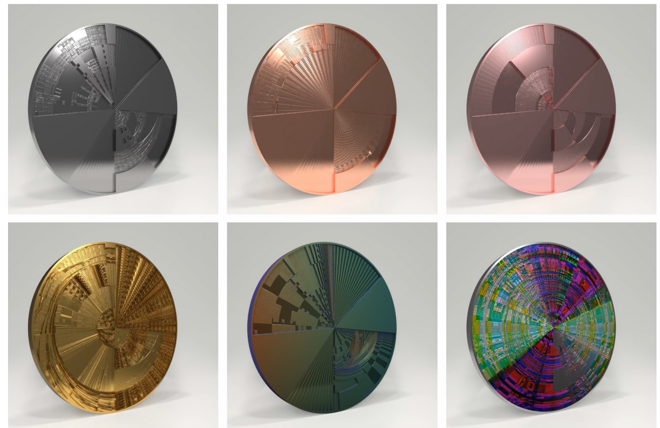

Figure 11: Six digital assets demonstrating the different materials used to achieve six rarity levels. From top left to bottom right: nickel, copper, rose gold, gold, chameleon, and RGB – special level with 40 digital assets.

7 Conclusions

The use of AI in creative and artistic contexts has garnered significant interest. Following on this line, INCM invited us to use AI to aid in designing a commemorative coin themed around the "digital world." Designing coins requires balancing practical functions and aesthetic qualities, considering factors such as mintability, durability, recognisability, legal requirements, thematic significance, and visual appeal.

Our goal with this project was to create a coin that honours the traditions of numismatics and the Portuguese National Press-Mint, while also reflecting the identity of our AI and Design Lab and our perspectives on Computational Creativity (CC) research. We aimed to symbolise the "digital world" and explore how co-creativity can enhance traditional numismatic art, bridging the gap between coin minting heritage and computational aesthetics. Additionally, the coin also references how humans and machines are increasingly inseparable, literally two faces of the same coin. It illustrates how human-machine cooperation and Human-Centred AI can expand human creative potential and foster ethical use that transcends mere imitation of Human artistry.

Acknowledgments

The authors would like to express their gratitude to Imprensa Nacional–Casa da Moeda for presenting us with the challenge that originated this research. Part of this work was developed during Penousal Machado's sabbatical stay at the BEACON Center for the Study of Evolution in Action at Michigan State University. We would like to thank Wolfgang Banzhaf, BEACON, and MSU for hosting and supporting this endeavour. Special thanks to Fábio Gouveia for his crucial role in generating the digital assets. This work is funded by the FCT - Foundation for Science and Technology, I.P./MCTES through national funds (PIDDAC), within the scope of CISUC R&D Unit - UIDB/00326/2020 or project code UIDP/00326/2020.

References

[Anwar *et al.*, 2021] Hafeez Anwar, Saeed Anwar, Sebastian Zambanini, and Fatih Porikli. Deep ancient roman republican coin classification via feature fusion and attention. *Pattern Recognition*, 114:107871, 2021.

[Baeta *et al.*, 2021] Francisco Baeta, João Correia, Tiago Martins, and Penousal Machado. TensorGP - genetic programming engine in TensorFlow. In *EvoApplications*, volume 12694 of *Lecture Notes in Computer Science*, pages 763–778. Springer, 2021.

[Banzhaf and Machado, 2024] Wolfgang Banzhaf and Penousal Machado. Fundamentals of evolutionary machine learning. In Wolfgang Banzhaf, Penousal Machado, and Mengjie Zhang, editors, *Handbook of Evolutionary Machine Learning*, pages 3–28. Springer Nature Singapore, Singapore, 2024.

[Berns and Colton, 2020] Sebastian Berns and Simon Colton. Bridging generative deep learning and computational creativity. In *ICCC*, pages 406–409. Association for Computational Creativity (ACC), 2020.

[Boden, 2004] Margaret Boden. *The Creative Mind: Myths and Mechanisms*. Routledge, 2004.

[Boden, 2009] Margaret A. Boden. Computer models of creativity. *AI Mag.*, 30(3):23–34, 2009.

[Brown and Jordanous, 2022] Daniel Brown and Anna Jordanous. Is style reproduction a computational creativity task? In *ICCC*, pages 220–229. Association for Computational Creativity (ACC), 2022.

[Cohen and McCorduck, 1995] Harold Cohen and Pamela McCorduck. *AARON's Code: Meta-Art, Artificial Intelligence, and the Work of Harold Cohen*. W. H. Freeman & Co, 1995.

[Colton and Wiggins, 2012] Simon Colton and Geraint A. Wiggins. Computational creativity: The final frontier? In *ECAI*, volume 242 of *Frontiers in Artificial Intelligence and Applications*, pages 21–26. IOS Press, 2012.

[Colton *et al.*, 2011] Simon Colton, John William Charnley, and Alison Pease. Computational creativity theory: The FACE and IDEA descriptive models. In *ICCC*, pages 90–95. computationalcreativity.net, 2011.

[Cope, 2005] David Cope. *Computer Models of Musical Creativity*. MIT Press, 2005.

[Correia *et al.*, 2019] João Correia, Penousal Machado, Juan Romero, Pedro Martins, and F. Amílcar Cardoso. Breaking the mould an evolutionary quest for innovation through style change. In *Computational Creativity*, pages 353–398. Springer, 2019.

[Correia *et al.*, 2024] João Correia, Francisco Baeta, and Tiago Martins. Evolutionary generative models. In Wolfgang Banzhaf, Penousal Machado, and Mengjie Zhang, editors, *Handbook of Evolutionary Machine Learning*, pages 283–329. Springer Nature Singapore, Singapore, 2024.

[Cunha *et al.*, 2019] João Miguel Cunha, Nuno Lourenço, João Correia, Pedro Martins, and Penousal Machado. Emojinating: Evolving emoji blends. In *EvoMUSART*, volume 11453 of *Lecture Notes in Computer Science*, pages 110–126. Springer, 2019.

[Davis *et al.*, 2016] Nicholas M. Davis, Chih-Pin Hsiao, Kunwar Yashraj Singh, Lisa Li, and Brian Magerko. Empirically studying participatory sense-making in abstract drawing with a co-creative cognitive agent. In *IUI*, pages 196–207. ACM, 2016.

[Du Sautoy, 2019] Marcus Du Sautoy. *The creativity code: Art and innovation in the age of AI*. Harvard University Press, 2019.

[Eiben and Smith, 2015] Ágoston E. Eiben and James E. Smith. From evolutionary computation to the evolution of things. *Nat.*, 521(7553):476–482, 2015.

[Elgammal *et al.*, 2017] Ahmed M. Elgammal, Bingchen Liu, Mohamed Elhoseiny, and Marian Mazzone. CAN: creative adversarial networks, generating "art" by learning about styles and deviating from style norms. In *ICCC*, pages 96–103. Association for Computational Creativity (ACC), 2017.

[Espírito Santo *et al.*, 2023] Luís Espírito Santo, André C Santos, and Marcio Lima Inácio. Focusing on artists' needs: Using a cultural probe for artist-centred creative software development. In *Proceedings of the Eighth International Conference on Computational Creativity, ICCC 2023*. Association for Computational Creativity (ACC), 2023.

[Goodfellow *et al.*, 2014] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron C. Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, pages 2672–2680, 2014.

[Hachman, 2022] Mark Hachman. Midjourney's enthralling AI art generator goes live for everyone, July 2022. Archived from the original on October 20, 2022. Retrieved September 5, 2022.

[Ho *et al.*, 2020] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.

[Jordanous, 2012] Anna Jordanous. A standardised procedure for evaluating creative systems: Computational creativity evaluation based on what it is to be creative. *Cogn. Comput.*, 4(3):246–279, 2012.

[Jordanous, 2017] Anna Jordanous. Co-creativity and perceptions of computational agents in co-creativity. In *ICCC*, pages 159–166. Association for Computational Creativity (ACC), 2017.

[Karras *et al.*, 2019] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, pages 4401–4410. Computer Vision Foundation / IEEE, 2019.

[Machado and Cardoso, 2002] Penousal Machado and Amílcar Cardoso. All the truth about NEvAr. *Appl. Intell.*, 16(2):101–118, 2002.

[Machado *et al.*, 2008] Penousal Machado, Juan Romero, and Bill Z. Manaris. Experiments in computational aesthetics. In *The Art of Artificial Evolution*, Natural Computing Series, pages 381–415. Springer, 2008.

[Machado *et al.*, 2014] Penousal Machado, Tiago Martins, Hugo Amaro, and Pedro Henriques Abreu. An interface for fitness function design. In *EvoMUSART*, volume 8601 of *Lecture Notes in Computer Science*, pages 13–25. Springer, 2014.

[Machado *et al.*, 2021] Penousal Machado, Juan Romero, and Gary Greenfield. *Artificial Intelligence and the Arts – Computational Creativity, Artistic Behavior, and Tools for Creatives*. Springer, 2021.

[Martins *et al.*, 2018] Tiago Martins, João Correia, Ernesto Costa, and Penousal Machado. Evotype: Towards the evolution of type stencils. In *EvoMUSART*, volume 10783 of *Lecture Notes in Computer Science*, pages 299–314. Springer, 2018.

[Martins *et al.*, 2023] Tiago Martins, João Miguel Cunha, João Correia, and Penousal Machado. Towards the evolution of prompts with metaprompter. In *EvoMUSART@EvoStar*, volume 13988 of *Lecture Notes in Computer Science*, pages 180–195. Springer, 2023.

[Oppenlaender, 2022] Jonas Oppenlaender. The creativity of text-to-image generation. In *MindTrek*, pages 192–202. ACM, 2022.

[Parente *et al.*, 2023] Jéssica Parente, Luís Gonçalo, Tiago Martins, João Miguel Cunha, João Bicker, and Penousal Machado. Using autoencoders to generate skeleton-based typography. In *EvoMUSART@EvoStar*, volume 13988 of *Lecture Notes in Computer Science*, pages 228–243. Springer, 2023.

[Pereira *et al.*, 2019] Fábio André Pereira, Tiago Martins, Sérgio Rebelo, and João Bicker. Generative type design: Creating glyphs from typographical skeletons. In *ARTECH*, pages 19:1–19:8. ACM, 2019.

[pharmapsychotic, 2023] pharmapsychotic. CLIP interrogator: Image to prompt with BLIP and CLIP. https://github.com/pharmapsychotic/clip-interrogator, 2023. Accessed: 2023-09-01.

[Pressburg Mint, 2023] Pressburg Mint. AI Coin 2023 1oz Silver 9999 Bullion. https://pressburgmint.com/product/ai-coin-2023-1oz-silver-9999-bullion/, 2023. Accessed: 2024-01-17.

[Radford *et al.*, 2021] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 2021.

[Ramesh *et al.*, 2021] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, volume 139 of *Proceedings of Machine Learning Research*, pages 8821–8831. PMLR, 2021.

[Rebelo *et al.*, 2024] Sérgio M. Rebelo, Juan Julián Merelo, João Bicker, and Penousal Machado. Evaluation metrics for automated typographic poster generation. In *EvoMUSART*, volume 14633 of *Lecture Notes in Computer Science*, pages 326–341. Springer, 2024.

[Rombach *et al.*, 2022] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, pages 10674–10685. IEEE, 2022.

[Romero and Machado, 2008] Juan Romero and Penousal Machado, editors. *The Art of Artificial Evolution: A Handbook on Evolutionary Art and Music*. Natural Computing Series. Springer, 2008.

[Schuhmann *et al.*, 2022] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5B: an open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022.

[Silva and Costa, 2009] Sara Silva and Ernesto Costa. Dynamic limits for bloat control in genetic programming and a review of past and current bloat theories. *Genet. Program. Evolvable Mach.*, 10(2):141–179, 2009.

[Sims, 1991] Karl Sims. Artificial evolution for computer graphics. In *SIGGRAPH*, pages 319–328. ACM, 1991.

[Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, pages 5998–6008, 2017.

[Veale and Cardoso, 2019] Tony Veale and F. Amílcar Cardoso, editors. *Computational Creativity - The Philosophy and Engineering of Autonomously Creative Systems*. Springer, 2019.

[Ventura, 2016] Dan Ventura. Mere generation: Essential barometer or dated concept? In *ICCC*, pages 17–24. Sony CSL Paris, France, 2016.

[Wiggins, 2006] Geraint A. Wiggins. A preliminary framework for description, analysis and comparison of creative systems. *Knowl. Based Syst.*, 19(7):449–458, 2006.

[Yannakakis *et al.*, 2014] Georgios N. Yannakakis, Antonios Liapis, and Constantine Alexopoulos. Mixed-initiative co-creativity. *Proceedings of the 5th International Conference on Computational Creativity*, pages 150–159, 2014.