

Font Your Friends and Loved Ones: On the Utility of Ugly Interfaces

Josh Urban Davis

Department of Computer Science, Dartmouth College
josh.u.davis.gr@dartmouth.edu

Johann Wentzel

School of Computer Science, University of Waterloo
jdwentze@uwaterloo.ca



Figure 1: typeFACE is a generative font, formed from the output of a generative adversarial network trained on human faces. typeFACE stands as our declaration of independence from the dogma of clean, hyper-optimized aesthetic interface design.

ABSTRACT

User interface design often focuses so heavily on clean and minimal interface aesthetics that any deviation is often rejected as “ugly”. This tendency towards abstraction in UI design can be contextualized as a removal of the “human” or “physical world” from the aesthetic choices and design considerations for the system. To resist this techno-deterministic eradication of the human presence from UI design, as well as radically inject the human presence back into user interfaces, we present typeFACE, a web interface and generative adversarial network designed to create fonts from human faces. We provide an implementation and applications for such a system, as well as contextualize and analyze the history of “ugliness” and the “uncanny” in UI design history. We also discuss implications of such a system within the domains of data ownership, identity, and HCI design research.

CCS CONCEPTS

• **Computing methodologies** → *Neural networks*; • **Human-centered computing** → **Interaction design theory, concepts and paradigms**; *User interface design*; **HCI theory, concepts and models**.

KEYWORDS

neural networks, typefaces, speculative futures

ACM Reference Format:

Josh Urban Davis and Johann Wentzel. 2021. Font Your Friends and Loved Ones: On the Utility of Ugly Interfaces. In *CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '21 Extended Abstracts)*, May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3411763.3450371>

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

CHI '21 Extended Abstracts, May 8–13, 2021, Yokohama, Japan

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8095-9/21/05.

<https://doi.org/10.1145/3411763.3450371>

1 INTRODUCTION

A common priority within the HCI and UX research communities is the development of “clean”, minimalist user interfaces [34]. Interfaces must often meet a specific balance of usability, accessibility, and aesthetic appeal. At present, typical user interfaces are flat, modern, and optimized for subtly instilling purchasing habits and addiction-level attachments to scrolling through content. The benefit of minimal and clean user interfaces is that they do not “obstruct” or hinder the user in any particular way while the user performs their task [33]. However, as the feature sets and common use cases of interfaces evolve, so too do their interface design philosophies. User interfaces in modern systems have followed several trends, including interface elements matching their real life counterparts with detailed textures [39], matching real physical behaviours common in prototyping [19], and accurate behaviour of light among several levels of depth [38].

In many cases, interfaces designed specifically for engagement, usability, and aesthetic cleanliness can lose an element of humanity. Everyday physical objects tend to have obvious imperfections because they were created by an imperfect human hand. Likewise, machines typically work with a level of precision unattainable by humans. As a result, imperfections often betray the presence of a human in the creation process. The aesthetic evoked by clean, minimal interfaces mimic the automation of labor by abstracting away the human presence in exchange for cleanliness and machinic perfection. While minimal cleanliness is an important factor in interface inclusivity and usability, the philosophy of aesthetics underlying this design movement has infiltrated our wider culture. In a digital landscape that increasingly focuses on shaving down the milliseconds between a user seeing a product and making a purchase, interface elements that take ugliness as a point of pride rather than a mark of shame or apathy can add an element of familiarity, or even trustworthiness to otherwise bland spic-and-span interfaces.

This effect is even pervasive in tech culture: social media users create interfaces of themselves [23] by carefully curating images and ideologies to associate with their physical person [20]. Our constant submersion in a culture curated to appear clean results in

an unconscious awareness of this artifice. As such, deviating from that pattern must signify a deviation from artifice. If we are aware that “clean” and “curated” things must be artificial, then things that are “ugly” must be “true” [21]. However, “ugliness” is no more authentic than “cleanliness” in interfaces, as within this context intentional ugliness is just as curated and artificial. It is because our culture standardizes minimal interfaces that the disruption of these clean surfaces is perceived as “truth” [23]. Even outside of design, appearing unfiltered or unrefined in the presence of an overly-polished norm can be perceived as being more truthful. This has several historical examples, such as within politics, social media, and online culture as a whole.

Clean interfaces not only remove the human presence from their surfaces, but pose the danger of priming cultural exploitation and negative epistemology [22]. Typography, arguably the first user interface [36], defines its user experience purely through its visual design. The primary functional factors of typography are legibility, while the user experience in typography is derived solely through the feelings evoked by its appearance. Minimalist aesthetics have been common throughout the history of typeface design, due to their efficiency and legibility. The font Helvetica, for example, was developed in the explicit interest of “a neutral typeface that should give no additional meaning” [36]. While Helvetica is still common, many designers often criticize this “true neutral” philosophy, citing that “no design is truly empty of ideology” [12]. Criticisms of Helvetica and its associated philosophy of minimalism became more common after Apple’s 2013 choice to elect Helvetica Neue as its default interface font, subsequently replacing it after only two years.

The goal of visual precision and minimalism is the implied removal of the human hand from creation, supplanting it with the perfection of machinery. We therefore propose an alternative future that embraces the “ugly” by embracing “the human”. To explore this goal, we describe an example of an intentionally ugly design called typeFACE, a system that can create a font out of the facial features of any user (Figure 1). typeFACE uses a generative adversarial network trained on images of a user’s face and a selection of various typefaces to create fonts from human faces. To understand the implications of such a system, we perform an analysis of the history of “ugliness” as a cultural aesthetic, and examine how this construct intersects with the history of UI design. Exploring and deconstructing the particular evocation of “ugliness” evoked by typeFACE reveals a dialogue between the “uncanny” or “situational” horror latent within such systems. To faithfully experience the uncanniness of typeFACE, we encourage readers to engage with our prototype web interface¹ where visitors can render a chosen text in typeFACE. We present this work as a radical and playful gesture towards injecting the human and imperfect back into interface design research, and encourage an embrace of the ugly as avant-garde UI design.

2 BACKGROUND

Our work is a generative neural network typeface that embraces the uncanny valley. We contextualize our work within the history of ugliness, UI/UX design, and new media interface practice.

¹gofontyourself.xyz

2.1 Who Are You Calling Ugly?

“Ugliness” in the west has historically comprised many meanings that shift over cultures and time. The classical world perceived “beauty” as the idealized and the perfect, and as such “ugliness” encompassed all things outside this definition of perfection: the imperfect [13]. The concept of ugliness can be thought of as a historical fascination with the imperfect. Exploring cultural artifacts which are deemed “ugly” can reveal qualities that a given society deems physically and morally undesirable. Therefore, what a society deems ugly serves as a reflection of this society’s perception of the universe’s imperfections. For example, artists of ancient Greece saw the universe as a flawed reflection of the perfected realm of ideas, and their statues of idealized pantheon bodies reflected this ideal [4]. Plotinus further defines ugliness as “flawed” or “evil”, strongly resembling the physical world [41].

The definition of beauty, and thus ugliness, is relative to the political culture evaluating it. The early Christians depicted Christ as tortured and humiliated to reflect the potential for redemption in all things earthly and divine, while the late Christians depicted the apocalypse, Satan, and death as morality horrors that plague the human experience [13]. Death is consistently portrayed as a counterpart to the devil throughout the Middle Ages, which is unsurprising given the ravages of the Black Death. The desecration and mutation of the body continues to haunt horrific depictions of morality tales throughout the Renaissance and into the 20th century. These examples are not intended to sketch a concise history of ugliness, but serve to illustrate how definitions of morality surrounding these aesthetics conform to the politics of that time. Ugliness, therefore, is a culturally-relative dialogue of aesthetics and ideas reflecting what is deemed “imperfect” by a given political structure at a specific time. Analysis of this aesthetic can not only reveal what a political structure sees as undesirable or imperfect, but also belies a persistent human obsession with the dark.

2.2 Ugly Design and User Interfaces

Similar to the analysis above, we can approximate an understanding of ugliness in user interfaces by examining definitions of perfection and their opposites in the history of UI design. Historians often point to the tradition of feng shui in 4000 BCE China as the origin of UI design, in which the arrangement of objects in a space was curated to promote well-being [58]. In this context, ugliness would consist of any arrangements of such objects which did not promote contentment and positivity. Similarly, in the 5th century BCE in the West, Hippocrates first described the preferred arrangement of surgeons’ quarters to allow quick and easy access to tools [26]. While many historical examples of preferable design have been documented, the notion of “bad design” is a relatively recent invention. One of the earliest written descriptions of bad design comes from Henry Dryfuss’ 1955 treatise “Designing for People” where he states “when the point of contact with the product and people becomes a point of friction, then the [design] has failed” [11]. Successful and unsuccessful designs are often couched in the language of utility as opposed to “ugly” or “beautiful”, yet the dichotomy between flawless utility and imperfect uselessness is still present [48]. A more direct association between design considerations and aesthetics is evident in the Aesthetic-Usability effect, in which users perceive

an aesthetically pleasing system as being more effective [37]. This reflects that cultural definitions of beauty and ugliness still impact the perceived usefulness of a system. In this way, the aesthetics help shape the utility of the interface, and not the other way around.

2.3 Generative Art and Design

While this dialogue of ugliness is evident in the history of art and design, little is present within the domain of UI research. Lev Manovich suggests that the new media art movement could be read as a response to emerging technology, providing radical alternatives to how interface technologies are designed [30].

One recent incarnation of new media arts uses algorithms as its primary medium. Generative art allows a designer to input high-level parameters describing their design goals, and algorithmically generates a solution space that satisfies those constraints. Several previous works examine generative art from the perspective of art theory [5, 10, 16], as well as generally categorize generative art as a midpoint in the spectrum between technical and cultural contribution [54]. Previous examinations of generative art discuss the “computational sublime” [31], posing that generative art is more than the algorithms and code that go into it, similarly to how traditional art forms are more than the sum of the tools used to create them. Previous technical contributions to generative art exist in previous literature, such as examples of generative artwork inspired by nature photos [9]. Our work draws from generative art history to imagine a future in which generative art can directly inform interface design.

Similarly, the intersection of neural networks and typefaces consists of two general topics: character recognition and font generation. Several previous tools have been developed to recognize characters and unique strokes. Works by Muriel Cooper [46] and Jürg Lehni [29] played foundational roles in establishing computational typefaces as a design medium. Including classic examples of stroke recognition [56], previous work also examines the use of neural networks to recognize characters in text in both English [42] and Chinese [59]. Similarly, Deepfont identifies fonts from images [53]. More directly related to our work is font generation, accomplished either through style transfer or through novel generation. Style transfer, the process in which a neural network is trained on previous fonts and images, applies these pre-trained styles to pre-existing fonts. Examples of style transfer include works by Atarsaikhan et al. [2] and Upchurch et al. [52]. Novel generation, which generally uses generative adversarial networks to create novel fonts from pre-existing images, has been accomplished in previous works [25, 47].

3 SMARTPHONES AND THE RE-DEFINITION OF AESTHETICS

The advent and eventual ubiquity of the smartphone prompted dramatic changes in the overall artistic direction of interfaces. Equipped with smaller touch-based screens and tasked with displaying rich, interactive content, mobile user interfaces emphasized usability and clarity. Being that smartphones are the most common device for interacting with online content [7], it makes sense

that smartphone-based UI trends propagate outward to other interface technologies. We describe some of these trends and how they pertain to modern UI design.

First came skeuomorphism, the practice of representing items in a design through depicting their real-life counterparts. This includes elements like realistic behaviour of light and shadow, as well as finer-grained realistic textures. As an example, early iterations of the iOS Notepad app used a font similar in look to physical writing, and the background of the app resembled a physical notebook. From a theoretical standpoint, one theory is that before iOS devices (and smartphones in general) became ubiquitous, designing interface elements similarly to their real-world counterparts pre-establishes feelings of trust and rapport with a device or application. Curtis [8] writes:

“Through this close mirroring in appearance and functionality, users are more readily able to relate iOS applications to their historical digital ancestors. This in turn may help the user establish a rapport with the skeuomorphic iOS interface. In this way iOS skeuomorphs act as visual metaphors, which are aided by being the product of the historical, cultural evolution of Mac OS”.

Next came the more recent implementation of flat designs. Instead of the complex shading, reflective surfaces, and detailed textures in many designs at the time, Apple’s major UI overhaul in iOS 7 sparked a widespread trend in digital interfaces toward flat design. Flat design, first popularized in Russia, the Netherlands, and Germany in the 1920s [57], focuses on readability, cleanliness, and minimalism. Within the purview of app interfaces, this means that most interfaces opted toward flatter color schemes with less depth. A more understated interface allows onscreen content to be the main point of focus.

More recently, two smaller trends have become common in web design and smartphone application design. Material design, first popularized by Android devices, renders all UI elements and panels on “cards”, with shadows providing the depth cues and hierarchy of content. A toned-down sibling to skeuomorphism, neumorphism accurately portrays the behaviour of light among several layers of application depth but keeps the simple, minimal textures of flat design.

4 THE AVANT-GARDE AND A RETURN OF THE HUMAN TO THE MACHINE

Skeuomorphism and flat design can be read as an aesthetic dialogue concerning what role the real world should serve in user interfaces; a dichotomy between representing tools through visual metaphors from the physical world and representing tools using smooth abstracted surfaces. However, this debate still designates utility and frictionless interaction as a primary motivator. These constraints limit the ability of designers and researchers to develop interfaces where friction could potentially be useful. In early Western music composition, composers were encouraged to pursue harmony and avoid dissonance in their work [13]. This tradition continued into the medieval and Renaissance eras, only fully being disassembled with the emergence of the avant-garde in the 20th century. Avant-garde artists and composers used dissonance to create tension, and harmony to relieve tension. Furthermore, literature in psychoacoustic research indicates that dissonant intervals can be used to

excite, and have been employed by composers since the thirteenth century to produce determined effects in appropriate contexts [20]. By reexamining the debate between skeuomorphism and flat design informed by this history, we can better explore their aesthetic effect. The "ugliness" and imperfection of the real world creates tension, while the relative cleanliness of an abstract design can relieve it.

Often, the avant-garde's role in cultural history was to fiercely experiment with offensive aesthetics in order to question and redirect our understanding of beauty [13]. Thus, a debate constrained by the unwavering desire for utility in UI design overlooks the potential for the useless to enrich our experience of interaction. Dunne and Raby highlight this concept in "Speculative Everything", encouraging playful engagement with speculative futures through design [12]. Similar to the avant-garde artists of the early 20th century, the goal of such work is to subtly nudge the direction of aesthetic development. Unlike the avant-garde who sought to challenge prevailing notions of beauty and ugliness in order to direct future aesthetic considerations, Dunne and Raby encouraged a similar practice of designing speculative futures in order to direct and change the development of "preferable futures" from "current futures". It is with this spirit that we developed typeFACE as an interrogative process to examine the impact of intentional ugliness within future HCI design research.

5 THE USES OF THE UNCANNY

In the early 20th century, Freud first described the uncanny as a "situational" ugliness, where mundane experiences are made horrific given the context in which they occur [15]. Ernst Jentsch went even further describing the uncanny as "intellectual uncertainty" and a psychological phenomenon reflecting something "we can't figure out" [27]. A common contemporary experience of the uncanny concerns the "uncanny valley effect", which refers to the discomfort viewers feel when a human facsimile closely resembles, but is not convincingly, a human. The situational ugliness here occurs because context implies that we are observing a human, yet the human facsimile (e.g. A 3D rendering of humans from *The Polar Express* [51]) is not believable. First described within human-robot interaction [35], the uncanny valley is a common point of reference when designing humanoid or human-like avatars.

Contemporary examples of fascination with the uncanny are evident in many micro-genres of internet art. The "liminal spaces" trend [32], for example, involves images of otherwise uninteresting rooms which become vaguely unsettling when divorced from their original context. Liminal spaces evoke a strange sense of dread since it is unclear why these typically nostalgic spaces are devoid of people and objects [51]. Similarly, the "#blursed" [44] trend on social media websites like Reddit and Twitter involves images experimenting with situational ugliness. #blursed evolved from a combination of the #blessed image trend (images depicting something pleasant or wonderful) [43] and the #cursed image trend (often depicting situations of abject horror within otherwise neutral contexts) [45]. Examples of #blursed often blend situational horror by recontextualizing #blessed images in a #cursed fashion. For example, a #blessed image may depict a photograph a fun photo with a cartoon character in costume, a #cursed image may depict a dark boiler room with a lurking character in a distressed and dirty



Figure 2: Input of human face and letter producing a single letter sample by minimizing the difference between the two images. The network was weighted heavily to prefer the face image in order to preserve as much of the facial structure as possible.

costume, and a #blursed image may depict an otherwise appealing character drawn in a strange or unsettling way. The uncanny horror in these examples is evoked by the situational recontextualization of familiar childhood characters into recognizable, but unfamiliar forms.

This situational horror highlights the imperfection evident in human interpretation and perception of childhood. The imperfection of limited human understanding evidenced by this aesthetic collision evokes a horror rooted in the uncanny and grotesque delicacies of the human animal. As Mark Fischer in "The Weird and the Eerie" writes, "We could go so far as to say that it is the human condition to be grotesque, since the human animal is the one that does not fit in, the freak of nature who has no place in the natural order and is capable of re-combining nature's products into hideous new forms" [14].

typeFACE sits within the realm of situational ugliness, intentionally evoking the uncanny by presenting a reader's text using characters derived from their face. The reader's awareness of this situation is what gives rise to typeFACE's emotional effect. Previous work in design describes strategies to mitigate this effect [17], but minimal previous work describes theoretical applications for intentionally evoking the uncanny. Our work observes this gap in the literature and provides a speculative look at ways to not only accept the uncanny valley effect, but embrace it for design utility.

6 TYPEFACE

We present typeFACE as a radical and playful gesture to re-infuse the personal, the ugly, and the human back into interface design. To do so, we used a Generative Adversarial Network (GAN) trained to minimize the difference between two image sets. One dataset contained a collection of images indicative of a given letter. This included 256 images of font letters per letter in the typeface. The second dataset contained a collection of 5 images taken from the user. GANs are appealing for this task because of their ability to persistently generate novel objects [9]. Goodfellow et al. proposed the

GAN which comprised two networks: a generator and a discriminator. The generator network synthesizes convincing objects to fool the discriminator [18]. Meanwhile, the discriminator attempts to distinguish between ground-truth objects and objects synthesized by the generator. The ground truth objects, in this case, consist of the two datasets of user face images and font letter images, with the system weighted heavily towards the user face images due to their scarcity. Training consists of the generator learning how to create images by adjusting weights corresponding to object features. Once trained, the resulting generator is able to produce, and interpolate between, images from the letter dataset and the input images of the user’s face. This follows an implementation similar to Neural Style Transfer and other image similarity approaches [2]. The resulting system is able to generate seemingly-infinite images for each letter that capture characteristics of both the user’s face and the legible letters (Figure 2). The benefit of using a GAN for this task is the variety of outputs a GAN is capable of generating. Most interpolation methods would fuse two images (the face and the target letter) which would generate a skin-covered version of the target letter. By using a GAN in combination with a dataset composed of many letters from many different fonts, we are able to generate unique typographic characteristics fused with the input face (e.g. serif placement, component thickness, etc). The result of this process is therefore a unique typeface, and not simply a skin-covered letter from another font.

7 FUTURE DEVELOPMENT IN THREE SMALL SPECULATIVE FICTIONS

Here we present three small speculative fictions in order to illustrate the conceptual impact of typeFACE on UI design. These playful examples are by-no-means exhaustive, but serve as creative speculation on the applications and ramifications of wider adoption of this design approach.

7.1 GO FONT YOURSELF DOT XYZ

To encourage playful interaction with uncanny UI design, we created an online interface² where users can input text to render it in typeFACE (Figure 3). We present this as a radical, playful gesture towards embracing the uncanny future this ugly approach presents, and encourage dialogue among the community regarding the experience of using such an interface. The current system allows a user to render text using a pre-trained font generating model, future exhibitions will allow user to upload their own facial data and create a font uniquely trained on themselves. We encourage readers to engage with our web interface system in order to properly experience the uncanniness evoked by reading their chosen text rendered in fonts derived from the human face.

7.2 THE FUTURE IS UGLY

Our current system outputs several images which must be manually sorted through to find the most suitable, flesh-like letters in the dataset. Future development of this work would automate this process. A semantic letter recognition system could be used to evaluate output letters for legibility, keeping only the most readable of

²<http://gofontyourself.xyz>

solutions. A further extension of this automation could similarly perform a facial recognition on the output letters, retraining the system to prioritize typefaces that are both legible and individually recognizable by facial recognition algorithms. This later development engenders a speculative fiction about the efficacy of facial recognition being extended into type. One can easily imagine a future where secure sign-and-seal lettering could be done using a person’s unique typeFACE. Furthermore, authorship integrity could be maintained by enforcing strict use of one’s own personal typeFACE for all documents. This speculative application could ensure that each digitally written word is traceable to its human source, at the expense of damaging online anonymity. The ethical teasing of this speculative fiction is worthy of further consideration when developing personalized interfaces.

7.3 LONG LIVE THE NEW FLESH

While we focused on fonts as the initial implementation of typeFACE, our vision extends beyond this to a future where fleshy, ugly, and disturbingly human aesthetics play a considerable role in UI design. While we contend that fonts are the first interface, one can easily imagine futures where all interfaces intersect with this proposed paradigm. Previous works within the human-factors community have explored ideas for such future interfaces including synthetic human-skin phone cases [50], robotic fingers for mobile devices [49], 3D-printed wearable algorithmically-generated representations of internal organs [40], and tangible data representations of human memory [1]. We see this work extending these predecessors by creating fleshy, typeFACE like interfaces for mobile phones, desktops, and other UIs. This future proposes a world where every interface a person interacts with is composed of their own bodily data. Phone interfaces can be mapped to a user’s face in a similar fashion as typeFACE, as could hyper-personalized desktops, or any other UI design interface that uses GUI and image based methods for interaction. Just as this approach can comprise the visual or aesthetic, so too could it be extended to the physical [55]. Similar to this, one can imagine an extreme implementation of typeFACE could comprise biological computers composed of users’ own personal data. These interfaces promise extreme personalization, with the added bonus that human biological material is compostable, and therefore more environmentally friendly than silicon.

8 DISCUSSION

Ugliness is useful as it forces us to reflect. typeFACE evokes a situational or uncanny ugliness by allowing a user to read a text they have written in a font rendered from their own face. The knowledge of this situation evokes a discomfort worthy of reflection. We discuss and explore the various “uglinesses” presented by our uncanny generating interface as a reflective exercise.

8.1 Data Horrors

Our system engages with the ongoing debate concerning the ethical ramifications of facial data. Several policies and legal debates cover the ownership and protection of personal data, which including facial image data. However, data willingly uploaded to the internet is often treated as an exploitable resource. Several facial recognition datasets are often composed of face images taken without prior



Figure 3: Example output of typeFACE, trained on a single user's facial data.

consent [3]. Visions of the future have been presented that entail users owning their own data, and being able to sell this data to various services if they choose [28]. However, this solution is not devoid of additional conflicts, as such a system could prey on the most vulnerable and economically disadvantaged. By rendering our written data using a type derived from our face, we are confronted with the question of who owns our data, and who should have access to such intimate information.

8.2 Identity Issues

Haraway's "A Cyborg Manifesto" playfully invokes the figure of the cyborg to explore how technologies are not ideologically neutral [24]. Since the invention of tools, humans have inexorably moved towards merging themselves with technology. However, as new technologies emerge, so too do opportunities for encoding such technologies with liberating as opposed to oppressive ideologies. typeFACE invokes the horror and responsibility of this revelation, as it visually presents a literal merger of human and machine to uncomfortable ends. This calls our attention to the ways in which technologies are developed, and how human identity is maintained against the backdrop of their powerful tools.

8.3 Meat in the Machine

The automation of human labor and experience presented by new technologies manifests a simple philosophy: the removal of the human presence [6]. Automation removes human laborers from the workforce and encourages the development of novel technologies to further this process. This is often rationalized using an intent to benefit, convenience, or help humans, and is often invoked as a panacea for most human problems. Similarly, we can see how removal of the human presence from user interfaces could invoke a similar techno-determinist philosophy, or at the very least, visually advocate and condone it. typeFACE invokes this horror by playfully returning the human to the system, to question and confront the removal of human presence from labor, interfaces, and futures.

9 CONCLUSION

Ugliness is relative to time and to culture. What was unacceptable yesterday may be acceptable tomorrow, and what is perceived as ugly may contribute, in a suitable context, to the beauty of the whole. Within the context of UI design and research, ugliness can be contextualized as a representation of the human or physical world within the design considerations of the system. To examine and confront this techno-determinist eradication of the human presence from UI design, we present typeFACE and encourage direct engagement with its particular ugliness. typeFACE participates in an aesthetic of ugliness we can contextualize as "uncanny" or "situational", being that the horror and discomfort arises from a

user reading a text written in a font derived from their own face. Reflecting on this discomfort reveals complications arising from data ownership and dissemination in a non-material world. Ugliness, as a result, serves a utility by inviting reflection and consideration of the horrors and discomforts coded into interactions. We encourage the adoption and further exploration of the latent potential within ugly interfaces for their ability to provoke, question, and contribute to our appreciation and understanding of human-computer interaction.

ACKNOWLEDGMENTS

Thanks to Justin Matejka (photos by John Laban) for volunteering his face to serve as training material for typeFACE, and to Fraser Anderson for support and encouragement during this ugly, beautiful project.

REFERENCES

- [1] Mirela Alistar and Margherita Pevere. 2020. Semina Aeternitatis: Using Bacteria for Tangible Interaction with Data. In *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems (CHI EA '20)*. Association for Computing Machinery, New York, NY, USA, 1–13. <https://doi.org/10.1145/3334480.3381817>
- [2] Gantugs Atarsaikhan, Brian Kenji Iwana, Atsushi Narusawa, Keiji Yanai, and Seiichi Uchida. 2017. Neural font style transfer. In *2017 14th IAPR International Conference on Document Analysis and Recognition (ICDAR)*, Vol. 5. IEEE, 51–56.
- [3] Ruha Benjamin. 2019. *Race After Technology: Abolitionist Tools for the New Jim Code*. John Wiley & Sons.
- [4] John Berger, Sven Blomberg, Chris Fox, Michael Dibb, and Richard Hollis. 1972. *Ways of Seeing*. British Broadcasting Corporation. Google-Books-ID: C4kBgwEACAAJ.
- [5] Margaret A Boden and Ernest A Edmonds. 2009. What is generative art? *Digital Creativity* 20, 1-2 (2009), 21–46.
- [6] Benjamin H. Bratton. 2015. *The Stack: On Software and Sovereignty*. MIT Press. Google-Books-ID: dEGdCwAAQBAJ.
- [7] J. Clement. 2020. Share of global mobile website traffic 2015-2020. <https://www.statista.com/statistics/277125/share-of-website-traffic-coming-from-mobile-devices/>. (Accessed on 12/14/2020).
- [8] Ambrose Curtis. 2015. Rhetoric of flat design and skeuomorphism in Apple's iOS graphical user interface. (2015).
- [9] Tom De Smedt, Ludvine Lechat, and Walter Daelemans. 2011. Generative art inspired by nature, using NodeBox. In *European Conference on the Applications of Evolutionary Computation*. Springer, 264–272.
- [10] Alan Dorin, Jonathan McCabe, Jon McCormack, Gordon Monro, and Mitchell Whitelaw. 2012. A framework for understanding generative art. *Digital Creativity* 23, 3-4 (2012), 239–259. <https://doi.org/10.1080/14626268.2012.709940> arXiv:<https://doi.org/10.1080/14626268.2012.709940>
- [11] Henry Dreyfuss. 2003. *Designing for People*. Allworth. Google-Books-ID: qOsJl7hW2qUC.
- [12] Anthony Dunne and Fiona Raby. 2013. *Speculative Everything: Design, Fiction, and Social Dreaming*. MIT Press. Google-Books-ID: 9gQyAgAAQBAJ.
- [13] Umberto Eco. 2011. *On Ugliness*. Maclehose. Google-Books-ID: G3lqkEACAAJ.
- [14] Mark Fisher. 2017. *The Weird and the Eerie*. Watkins Media Limited.
- [15] Sigmund Freud. 2018. *Das Unheimliche*. CreateSpace Independent Publishing Platform. Google-Books-ID: RXoHuwEACAAJ.
- [16] Philip Galanter. 2016. *Generative Art Theory: Paul/A Companion to Digital Art*. 146–180. <https://doi.org/10.1002/9781118475249.ch5>
- [17] Tom Geller. 2008. Overcoming the uncanny valley. *IEEE computer graphics and applications* 28, 4 (2008), 11–17.
- [18] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative

- Adversarial Networks. *arXiv:1406.2661 [cs, stat]* (June 2014). <http://arxiv.org/abs/1406.2661> arXiv: 1406.2661.
- [19] Google. [n.d.]. Design Guidelines - Material Design. <https://material.io/design>. (Accessed on 12/14/2020).
- [20] Boris Groys. 2014. *On the New*. Verso Books. Google-Books-ID: gslNBAAQBAJ.
- [21] Boris Grois. 2010. *Going Public*. Sternberg Press. Google-Books-ID: WvySZwEACAAJ.
- [22] Byung-Chul Han. 2017. *In the Swarm: Digital Prospects*. MIT Press. Google-Books-ID: JVaZDgAAQBAJ.
- [23] Byung-Chul Han. 2017. *Psychopolitics: Neoliberalism and New Technologies of Power*. Verso Books. Google-Books-ID: eSMcDgAAQBAJ.
- [24] Donna J. Haraway. 2016. *Manifestly Haraway*. U of Minnesota Press. Google-Books-ID: 4zB0DwAAQBAJ.
- [25] Hideaki Hayashi, Kohtaro Abe, and Seiichi Uchida. 2019. GlyphGAN: Style-consistent font generation based on generative adversarial networks. *Knowledge-Based Systems* 186 (2019), 104927.
- [26] Hippocrates. 2005. *Hippocratic Writings*. Penguin UK.
- [27] Ernst Jentsch. 1997. On the psychology of the uncanny (1906). *Angelaki* 2, 1 (Jan. 1997), 7–16. <https://doi.org/10.1080/09697259708571910>
- [28] Jaron Lanier. 2010. *You are Not a Gadget: A Manifesto*. Alfred A. Knopf. Google-Books-ID: u4jx5bSFlyEC.
- [29] Jürg Lehm. 2009. Typeface as Programme. *Typotheque* (2009).
- [30] Lev Manovich. 2002. Avant-garde as Software. *Artnodes* (Dec. 2002). <https://doi.org/10.7238/a.v0i2.681>
- [31] Jon McCormack and Alan Dorin. 2001. Art, emergence and the computational sublime. In *Proceedings of Second Iteration: A Conference on Generative Systems in the Electronic Arts*. Melbourne: CEMA, 67–81.
- [32] Know Your Meme. 2018. *Images With Elegiac Auras / Liminal Spaces*. <https://knowyourmeme.com/memes/images-with-elegiac-auras-liminal-spaces> (Accessed on 02/10/2021).
- [33] Per Mollerup. 2015. *Data Design: Visualising Quantities, Locations, Connections*. Bloomsbury Academic. Google-Books-ID: QvspBgAAQBAJ.
- [34] Mike Monteiro. 2019. *Ruined by Design: How Designers Destroyed the World, and what We Can Do to Fix it*. Mule Design. Google-Books-ID: haS0wwEACAAJ.
- [35] Masahiro Mori. 1970. The uncanny valley. *Energy* 7, 4 (1970), 33–35.
- [36] Lars Müller. 2005. *Helvetica: Homage to a Typeface*. Springer Science & Business Media. Google-Books-ID: 6UHOApOVw2EC.
- [37] Don Norman. 2013. *The Design of Everyday Things: Revised and Expanded Edition*. Basic Books. Google-Books-ID: nVQPAAAQBAJ.
- [38] David Ofiari. 2020. Neumorphism the right way — A 2020 Design Trend. <https://medium.com/@artofofiari/neumorphism-the-right-way-a-2020-design-trend-386e6a09040a>. (Accessed on 12/14/2020).
- [39] Tom Page. 2014. Skeuomorphism or flat design: future directions in mobile device User Interface (UI) design education. *International Journal of Mobile Learning and Organisation* 8, 2 (2014), 130–142.
- [40] William Patrick, Christoph Bader, Dominik Kolb, Sunanda Sharma, Steven Keating, and Neri Oxman. 2020. Mushtari. <https://mediatedmattergroup.com/mushtari>.
- [41] Plotinus. 1984. *Plotinus: Enneads*. Harvard University Press. Google-Books-ID: 1b_XzQEACAAJ.
- [42] A Rajavelu, Mohamad T Musavi, and Mukul Vassant Shirvaikar. 1989. A neural network approach to character recognition. *Neural networks* 2, 5 (1989), 387–393.
- [43] Reddit. [n.d.]. *top scoring links : blessedimages*. <https://www.reddit.com/r/blessedimages/top/> (Accessed on 02/10/2021).
- [44] Reddit. [n.d.]. *top scoring links : blursedimages*. <https://www.reddit.com/r/blursedimages/top/> (Accessed on 02/10/2021).
- [45] Reddit. [n.d.]. *top scoring links : cursedimages*. <https://www.reddit.com/r/cursedimages/top/> (Accessed on 02/10/2021).
- [46] David Reinfurt and Robert Wiesenberger. 2017. Muriel Cooper. *MIT Press* (Sep 2017).
- [47] Shusen Tang, Zeqing Xia, Zhouhui Lian, Yingmin Tang, and Jianguo Xiao. 2019. FontRNN: Generating Large-scale Chinese Fonts via Recurrent Neural Network. In *Computer Graphics Forum*, Vol. 38. Wiley Online Library, 567–577.
- [48] Frederick Winslow Taylor. 1913. *The Principles of Scientific Management*. Harper. Google-Books-ID: HoJMAAAAYAAJ.
- [49] Marc Teyssier, Gilles Bailly, Catherine Pelachaud, and Eric Lecolinet. 2018. Mo-biLimb: Augmenting Mobile Devices with a Robotic Limb. In *Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology*. 53–63.
- [50] Marc Teyssier, Gilles Bailly, Catherine Pelachaud, Eric Lecolinet, Andrew Conn, and Anne Roudaut. 2019. Skin-on interfaces: A bio-driven approach for artificial skin design to cover interactive devices. In *Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology*. 307–322.
- [51] Angela Tinwell. 2014. *The Uncanny Valley in Games and Animation*. CRC Press. Google-Books-ID: 8mZYBQAAQBAJ.
- [52] Paul Upchurch, Noah Snaveley, and Kavita Bala. 2016. From A to Z: supervised transfer of style and content using deep neural network generators. *arXiv preprint arXiv:1603.02003* (2016).
- [53] Zhangyang Wang, Jianchao Yang, Hailin Jin, Eli Shechtman, Aseem Agarwala, Jonathan Brandt, and Thomas S Huang. 2015. Deepfont: Identify your font from an image. In *Proceedings of the 23rd ACM international conference on Multimedia*. 451–459.
- [54] Mitchell Whitelaw. 2005. System stories and model worlds: A critical approach to generative art. *Readme* 100 (2005), 135–154.
- [55] Stephen Wilson. 2002. *Information Arts: Intersections of Art, Science, and Technology*. MIT Press. Google-Books-ID: sHuXQtYrNPYC.
- [56] Jacob O Wobbrock, Andrew D Wilson, and Yang Li. 2007. Gestures without libraries, toolkits or training: a \$1 recognizer for user interface prototypes. In *Proceedings of the 20th annual ACM symposium on User interface software and technology*. 159–168.
- [57] workerbee. 2017. Skeuomorphism vs. flat design vs material design - 99designs. <https://99designs.ca/blog/trends/skeuomorphism-flat-design-material-design/>. (Accessed on 12/14/2020).
- [58] Indi Young. 2008. *Mental Models: Aligning Design Strategy with Human Behavior*. Rosenfeld Media. Google-Books-ID: b5aLQ_cbC4kC.
- [59] Zhuoyao Zhong, Lianwen Jin, and Ziyong Feng. 2015. Multi-font printed Chinese character recognition using multi-pooling convolutional neural network. In *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*. IEEE, 96–100.