



A composite analysis of eyes and mouths of young Chinese women, created by DALL-E, Gemini, and Midjourney in a zero-shot setting.

Facial AI: Cosmetic Surgery in China and the Death of Internets Everywhere

MARGARET HILLENBRAND, XUENAN CAO

The Chinese internet is not the same as other internets, or so we are often told. Google, Facebook, YouTube, and X are all blocked, making for a palpably different online world, and in the absence of these behemoths local tech companies have filled the space with an energetic ecology of homegrown platforms. At the core of this ecology are China's distinctive cultures of celebrity and, in particular, the web of influencer-driven marketing via social media that shapes the Chinese digital economy and is emblemized by the term *wanghong*. This term literally means “net red,” or “hot online,” and refers in large part to the sway and seductive power of influencers, the mostly female internet celebrities who function as the lightning rod converting sexualized user attention into hard platform profit. At the core of their erotic capital is a highly bespoke look, a facial template that had become so entrenched by the late 2010s that it generated its own moniker: *wanghonglian*, which loosely translates as “Internet celebrity face.” The term rapidly gained intense media traction, spawning its own dedicated page on Baidu Baike, China's version of Wikipedia, in 2017. The site's definition of the look reads as follows:

Internet celebrity face . . . refers to a standardized face created through plastic surgery or makeup techniques, with features such as a pointed chin, European double eyelids, and a high-bridged nose.¹

Internet celebrity face soon achieved close to full saturation on China's money-spinning livestreaming platforms, whose audiences are predominantly male: as Yang Yidan notes, “93%, 96%, and 87% respectively of the 100 top-rated female hosts who get the most clicks have ‘Internet celebrity face.’”² The look is prevalent on the streets of urban China, too. Costly surgical procedures and cheap digital face-editing apps have made *wanghong* aesthetics attainable for any body or budget.

As a purchasable look that can be instantly capitalized in fate-changing ways, Internet celebrity face is controversial. There are screeds of posts on the Chinese net trolling women who have this look: haters claim that possessors of Internet celebrity face are identikit plastic cyborgs put together on an assembly line who have traded their authenticity for the chance to make a quick, lazy, mostly untalented buck. This invective is whetted

by the grim realities of China's bottlenecked job market, especially for young people. Why study for a degree/work in a factory/learn a trade when you can get a cosmetic surgery microloan to purchase an Internet celebrity face and make tens of thousands of *yuan* a day peddling products on a Douyin livestream? Debate on this topic reached a fever pitch in 2016, when Internet celebrity face became the "the most discussed buzzword of the year," in part because China's evolution into an increasingly lookist society is seen by many as a bellwether of mounting social injustice.³ Unsurprisingly, then, there has been significant pushback against Internet celebrity face over the last few years, as well as a concomitant social trend toward customization, personalization, and individuality in the domain of aesthetic



personhood. And the core driver for this apparent code-switch to all things idiosyncratic is artificial intelligence—a move that might seem counterintuitive at a time when GenAI is coming under suspicion as a force that flattens diversity.

Since 2020, the pathways to surgical aesthetic modification in China have become ever more entangled with AI and algorithmic governance—from cosmetics to surgery to in-app filters—all accompanied by loud marketing noise about the primacy of data-driven, made-to-order, personally tailored approaches. AI technology is permeating brick-and-mortar beauty retail sites in Chinese cities via so-called digital touchpoints that offer smart diagnostics and skin analysis powered by big data.⁴ Small wonder that French multinational beauty retailer Sephora recently chose China as one of its early test sites for in-store product recommendation based on AI-generated cosmetic trends. Cosmetic surgery in China is meshing with the field

of artificial intelligence even more comprehensively: from smartphone-based AI tools that evaluate patients' suitability for cosmetic surgery to large language models as facilitators of dialogue in the consultation process; and from the use of AI to enhance preoperative nasal contour design to AI-driven assessment of acid fillers as a treatment for boosting lip volume, all for the benefit of the individual—and individualized—patient.⁵ Wu Xinhong, the CEO of China's premier beauty filter app, Meitu, admitted in a 2021 interview with *Daily Economic News* (*Meiri jingji xinwen* 每日经济新闻) that in the past the company's products had “unified everyone's aesthetics” (via) so-called Internet celebrity face. But nowadays, he continued, “more and more flowers are blooming, and aesthetics are becoming more open and inclusive. We're gradually making everyone understand that their own beauty is unique.”⁶

This sounds like a laudable aim. But it is also another way of saying that to be idiosyncratic increasingly means editing oneself into so-called difference, a pursuit linked to facial modifications both surgical and digital. This latter practice is one that Meitu has been advancing aggressively for some years now. In the 2000s, beautifying a digitized image in its pixel space typically required hard-to-access professional software; but in October 2008 Meitu launched a tool named Meitu Xiuxiu 美图秀秀 (later known as MeituPic) with photo-editing heuristics written into it. On MeituPic, anyone could add filter effects and whiten the skin with one tap: the software made old-school Hollywood glamour available via a single button within a foolproof interface. Back in 2008, software of this kind was not known as AI but was instead understood as a set of rule-based image-processing algorithms.⁷ Like so-called traditional AI, this tool could edit a face because it was rule-based—that is, visual effects had to be prewritten into the program, with the various limitations that implies.⁸ To expand its features, the company later developed Meitu Mobile, an app with additional functions that could slim faces, narrow jaws, and enlarge eyes, all capabilities more effectively achieved through machine learning from data (rather than just prewritten filters).⁹

At first glance, digital facial modification might not look like artificial intelligence. AI, and generative AI as we know it today, was redefined by OpenAI's household technology ChatGPT in the winter of 2022. But by the 2010s, face-editing algorithms and generative models were already coproducing plenty of commercial applications. Indeed, some pre-GPT generative models developed in direct tandem with commercial applications of facial manipulation. A telling example is a face-swap app of 2012 called FaceFlip, which combines simpler computer vision techniques with image generators.¹⁰ Before the advent of this app, if users wanted to swap one face for another, they had to use two images of similar size and composition, create mask layers in an image-editing tool, and manually cover one face with another. By the mid-2010s, an appetite had emerged for apps that could automate this clunky manual process using

A transformative collage of images sourced from Xiaohongshu via the keyword “wanghonglian.” Original images were cropped to isolate the facial features, composited to highlight aesthetic homogeneity. Grayscale conversion was applied to emphasize formal similarities and minimize personal identification.

sophisticated computer vision techniques. The ready-made market for commercial-grade face manipulation tools has become a rich seedbed for the subsequent integration of generative models into face editing.

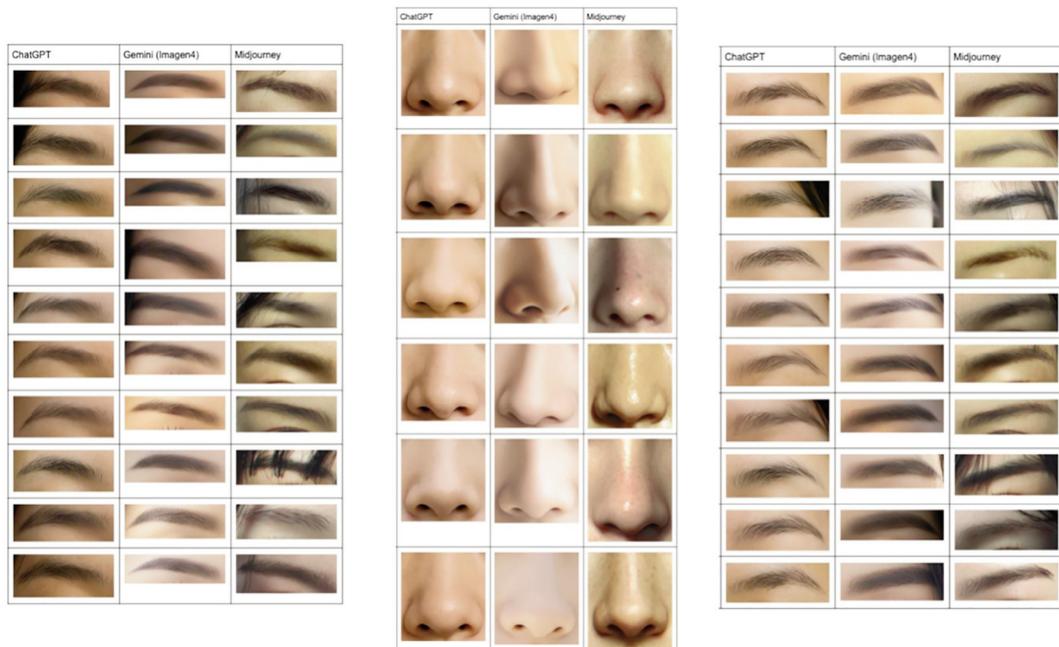
These AI-infused beauty cams and photo editors, which began to swarm the global app market, all claimed to provide professional-level edits at the touch of a fingertip. Users could now purchase the ability to lengthen the lashes, round out the forehead, and sharpen the jawline, all within the interface of a single app. Some tools promised to deliver individualized aesthetic personhood for a small subscription fee; others were advertised as “free.” On May 16, 2017, Facebook jumped on the bandwagon, introducing “an easy way to turn an ordinary selfie into something fun and entertaining,” which became an instant hit for users looking to upgrade their individuality. “Simply open the camera and tap the new face icon in the bottom right corner”: idiosyncrasy was right there for the taking.¹¹ Next, TikTok entered the competition with mature editing tools developed for the Chinese domestic internet; in 2018, it competed on the global market with its real-time editing tools and came out on top.

Within a few months, the war had begun on deepfakes, so named because these synthetic media use deep learning in order to insert celebrity faces into fake hardcore videos—and those of many more ordinary women into revenge porn and other forms of image-based sexual abuse. The pornographic deepfake was arguably the first truly viral use of generative models, exploding on forums that thrived on nonconsensual image-sharing and harassment. Over time, a broader deepfake market has emerged in which those without the requisite technical skills can simply commission someone else to create custom porn for as little as \$30 at time, thus escalating the age-old misogyny of the porn industry to novel techno-dystopian levels.¹² Female faces (and occasional male ones) are subject to nonstop, nonconsensual manipulations—deprecations achievable by digital techniques developed in tandem with those for bettering self-presentation. Harassment abounds, and banning it on social media has only limited effects.¹³ Indeed, Facebook and Bytedance, energetic coprogenitors of these facial AI tools, later acknowledged their infrastructural limitations in managing the ethical crises their own tech had spawned.

In short, facial editing and generative AI worked in collaboration to drive the worldwide commercial development of both voluntary facial editing and nonconsensual manipulation, utterly keystone trends that have in their turn migrated beyond specific apps to furnish the entire social media and livestreaming ecology with tools for both instant aesthetic modification and industrial-level sexual harassment. Nowadays, the inundation of the internet with AI-generated faces is cause for global consternation: from angsty news headlines to congressional hearings to belated and possibly toothless legislation.¹⁴ Back in the 2010s, though, the sentiment was quite different: at that

A composite analysis of noses and eyebrows from AI-generated images of young Chinese women, created by DALL-E, Gemini, and Midjourney in a zero-shot setting.

point, more AI-based tools meant more idiosyncrasy, more diversity, more user control over aesthetic personhood. Yet more crucial, and so far unremarked upon, is the fact that machine learning and the emergence of a huge market for face-editing tools that increasingly militate toward rigid homogeneity have not moved along separate paths as platform culture has become ever more awash with synthetic faces. To the contrary, the two have developed in tight symbiosis. Or to put this another way, the accelerating AI takeover of social media is much more about making and faking faces than has been acknowledged: without facial editing as spur and facilitator, the internet could never have become as comprehensively commandeered by AI as it now is. Our faces—traditionally the very seat of our humanity, the bedrock of who we are as people—have morphed into malware, into the most perfidious of Trojan horses.



The implications of this spiral when we return to the accelerating uniformity of facial editing: its drive to make faces conform to fixed beauty standards. When Meitu CEO Wu Xihong opines that “We’re gradually making everyone understand that their own beauty is unique,” he leaves much unsaid about what enabled these apparently “unique” services in cosmetic surgery and filter apps to emerge. Data analytics and AI-infused diagnostic, surgical, and editing procedures create an auto-catalytic loop. If you ask an image generator to produce ten images of young Chinese women, the outputs are disturbingly similar: the same perky nose, thick trimmed brows, large eyes with long lashes, and glossed rosebud lips. As an experiment, we prompted three generative models, OpenAI’s DALL·E, Google’s Gemini, and Midjourney, to output ten images of young Chinese women.¹⁵ Then, we compiled the features into six tables, each containing only the noses, eyes, eyebrows, or mouths of

the various generated images. The results, in the frontispiece and below, disclose the eerie similarity of these discrete features plucked from AI-generated Chinese faces. All the noses, eyes, eyebrows, and mouths, regardless of the generative model, obey the diktats of the Internet celebrity face template.

These results reveal that *wanghong* aesthetics are being generated by AI, pixel by pixel, as decisively as they are carved out by the scapel in operating theaters. As such, this *wanghong* look is an eerie revisitation of Walter Benjamin's famous analogy between the surgeon and the cameraman: "Magician and surgeon act like painter and cameraman," he writes. "The painter, while working, observes a natural distance from the subject; the cameraman, on the other hand, penetrates deep into the subject's tissue."¹⁶ Benjamin's notion of cinematographical invasiveness turns on the camera's ability to slice, segment, and dissect the face via close-up, angle, and lighting, to delve deep into the facial integrity of the filmic subject. Yet, if anything, the six tables above may bear an even closer genealogical relationship to the ideas of nineteenth-century godfather of eugenics Francis Galton, whose work overlaid—quite literally—techniques of photographic manipulation with statistical regression and pushed at that same boundary between the generic and the individual that is so germane to Internet celebrity face.¹⁷ Galton's composite photography, just like AI-generated Internet celebrity face, produced fixed facial categories (e.g., the "Jewish type," the "criminal type") from actually existing component features, forging the "universal" from granular splinters of individual faces.

Which is to say, the identikit faces that currently deluge the internet take their place within a lengthy lineage of photographic manipulation, in China as elsewhere. But whereas once upon a time Chinese artists of the doctored image worked in secret studios in the bowels of the Forbidden City, retouching photographs of Chairman Mao with tiny paintbrushes to give him an eternally youthful glow, nowadays image manipulation takes place in the wide open spaces of the internet using an incomparably more potent suite of tools.¹⁸ Yet this collusion between AI and Internet celebrity face has gone largely unnoticed even as the scaling of facial AI outputs—the trillions of faces that now populate social media and streaming platforms—is fed into the training of transformer-based foundation models, such as OpenAI's language and image models and Google's multimodal models. Social media imagescapes—already glutted with AI-edited aesthetics—have been systematically scraped for training data, as data-hungry companies competing in the race for the best-performing models devour images wherever they can find them. Given regulatory delays in data privacy and data protection, this gobbling has largely gone unchecked. In the context of the symbiotic development of facial editing tools and machine learning, it becomes clear that internet image content has long been contaminated by AI-edited faces, if not comprehensively dominated by such material. It should

thus come as no surprise that commercial image generators trained on online image content only know how to generate a Chinese countenance in the mold of Internet celebrity face.

In other words, the drive toward individuality in aesthetic personhood that we discussed earlier is an entirely disingenuous move, a fairly obvious smokescreen for its polar opposite. In some ways, this current tension between purported personhood and actual homogeneity recalls the debates between Kathy Davis and Susan Bordo on feminism and plastic surgery during the 1990s, in which Davis argued that women who get “work” done are not “cultural dopes” but rather active agents who choose surgery because of the economic/erotic/emotional benefits it can bring, and Bordo argued back that these procedures can scarcely be called elective or expressive of agency when the acceptable surgical palette for, say, a nose job is limited in the extreme: neat, straight, small.¹⁹ Or as Kathryn Pauly Morgan puts it, “More often than not, what appear at first glance to be instances of choice turn out to be instances of conformity. . . . As we more and more profoundly artifactualize our own bodies, we become more sophisticated archaeological repositories and records that both signify and symbolize our culture.”²⁰ Morgan’s use of the term “archaeological” here is perhaps distracting, since it is, above all else, a set of cutting-edge technological affordances that determine choice, now more than ever. Of course, rhinoplasty has always been a *techne*. But the migration of that *techne* and all its associated apparatus into one-stop beauty mega-apps, which provide information on clinics and surgeons, set up appointments, assist with microloans, and scan selfies so that algorithms can suggest tailored treatments, and which users navigate via the cell phones that are grafted onto their bodies like prosthetic limbs—and are thus as inescapable as they are convenient—makes choice seem ever more illusory.²¹

Just as Galton’s composite portraits generated spurious facial “norms,” so too does the recursive nexus between the digital and the surgical produce a pseudostandardized metric. But this time around, the target shifts from a two-dimensional racialized/criminalized image to the three-dimensional female face, and from a method that sought to identify types to one that actually realizes them both in the flesh and in every corner of the Chinese web—thus mounting a still more sustained and invasive assault on the age-old notion that personhood, in all its integrity, is lodged in the face. This pseudostandardized metric, a kind of Galtonian averaging updated for the era of facial manipulation technologies, turns the pursuit of distinctive aesthetic selfhood into a self-defeating quest. The vast diversity of human faces is flattened so as to assimilate the most liked—and, in the influencer economy, the most profitable—*wanghong* features. The coveted V-line jaw, for example, requires the insertion of artificial bone beneath chin tissues, or the use of an algorithm to change the pixels around the imaged mandible: scoring well on the jawline-to-cheekbone ratio means abiding

by the pseudostandardized metric for measuring influencers' monetizable look.²² And the logic of this metric decrees that distinctive facial features signal not individuality or personality but deviations from a tried-and-true formula.

All the while, these same faces are being pumped into China's internet ecology through advertisements and other marketing materials. On Taobao (the Chinese counterpart to Amazon), shop owners who once used image-editing tools to create basic product images can now generate more sophisticated advertisement images and even short videos using AI prompts and samples.²³ Blog posts about how to use generative AI in social media and e-commerce have soared in popularity, enabling influencers and shop owners to cut costs. Instead of signing actual flesh-and-blood influencers, who often come dogged by personal scandal, companies can create their own tailor-made, clean-cut AI influencers—mostly female characters who, unsurprisingly, possess Internet celebrity face. Chinese consumers in the 2020s are becoming increasingly habituated to the spectacle of AI-generated public figures, not just across e-commerce sites, but on official news channels and at sports events.²⁴ AI-generated news anchors, for example, can deliver preset content flawlessly and are deployed in high-profile, high-octane public events such as the Winter Olympic Games.²⁵ Living, breathing news anchors and influencers now have the power to extend their nationwide reach as their AI duplicates are hired for online local events, equipped with enhanced interactive features that enable them to customize their dealings with fans in charismatic ways.²⁶

This entrenchment of *wanghong* facial aesthetics via GenAI is more than disquieting on several fronts. In the most obvious sense, this standardization of beauty ideals has a dark undertow of misogyny, one which mirrors the findings of a 2024 UNESCO report on regressive gender stereotyping in large language models (LLMs). The report sets out to measure gender diversity in the content of AI-generated texts and found that open-source (and thus the most widely used) LLMs

tended to assign more diverse, high-status jobs to men, such as engineer, teacher, and doctor, while frequently relegating women to roles that are traditionally undervalued or socially stigmatized, such as “domestic servant,” “cook,” and “prostitute” . . . Llama 2-generated stories . . . about women made most frequent use of the words “garden,” “love,” “felt,” “gentle,” “hair,” and “husband.”²⁷

This textual embedding of gender bias within LLMs is already obvious cause for alarm. Yet it could be argued that the *visualization* of prescriptive female personhood via AI image generation constitutes an even more troubling step change, because the parameters for diversity narrow suffocatingly as women are now also expected to assume identical facial traits, via grueling surgical operations, time-consuming cosmetic labor, invasive filters, or outright substitution via bots.

Rather revealingly in this regard, the UNESCO report also noted that women constitute only 20 percent of employees in technical roles in leading global machine learning companies, a disparity closely matched by the mere 18 percent of authors at leading AI conferences who are female. There is, of course, nothing new here. Misogyny already blighted computer science during World War II and the Cold War, when female operators and programmers were employed on vital national projects but their skills and ingenuity were underpaid, underestimated, and even designated as “non-expert.”²⁸ This misogynistic trend has since hardened still further, particularly as unelected tech lords have stepped up their muscle flexes in the sphere of state governance. Only a month after the release of the UNESCO report, tech CEO Elon Musk sought to replace Diversity, Equity and Inclusion (DEI) with Merit, Excellence, and Intelligence (MEI), a barely coded macho motto that makes *diversity* a dirty word by equating it snidely with the lack of merit. In China, although tech leadership roles are largely reserved for men, female tech workers make up 46 percent of the total workforce.²⁹ This figure might seem at first sight to suggest a more optimistic future for tech women in China, but the AI-saturated imagescape continues to generate a waking nightmare on incessant repeat for young Chinese women. Or to put this another way, just as men disproportionately dominate the viewership of China’s *wanghong*-hosted livestreams, so are men massively over represented in the global AI industrial-intellectual complex through which Internet celebrity face has swamped online spaces and been propagated to millions of ordinary Chinese women.

Furthermore, this dissemination of prescriptive facial prototypes needs to be understood within the steadily dehumanized ecology of the Chinese internet—and this is a lesson for internets everywhere. Researchers on digital China are now finding that there is no need to be online, human, or even an AI bot to qualify as an internet celebrity. Once upon a time, the term denoted flesh-and-blood influencers; but it is now evolving into a “logic of everyday life” as the sharing of user-generated images and videos about physical places has set up new feedback loops between social media sites, urban locales, and attention as a marketable commodity.³⁰ *Wanghong* has become a kind of wraparound descriptor, applied to “a great many objects, contexts, aesthetics, spaces, and design practices in everyday life . . . [e.g.] *wanghong* city, *wanghong* space, *wanghong* beverage, *wanghong* architecture, *wanghong* gallery, and *wanghong* café.”³¹ Still more to the point, being *wanghong* is uncannily recursive. Initially, an urban site, space, or infrastructure might be designated *wanghong* because it is special; but over time it has become harder to tell what, exactly, makes a *wanghong* “thing” worthy of the name, as the descriptor sometimes gets attached to banal things and boring places to promote some kind of “glow up.” In this sense, to be *wanghong* is to partake in a logic of rigid but amorphous AI replication; its defining

characteristic is, paradoxically, a pervasive drag-and-drop sameness in which, to paraphrase Gertrude Stein, “there is (not much) there there.”³²

This recursivity both mirrors the use of Internet celebrity face in training GenAI and the integration of trained GenAI models in platforms propagating internet celebrity cultures. Zhang et al. note in their analysis of *wanghong* urbanism that the flattening of city space into *wanghong* digital spectacle “reproduces extant spatial and social inequalities.”³³ In much the same way, the entrenchment of facial prototypes via the alliance between face-editing and GenAI enforces cultural homogenization and works to evaporate real-life diversity within a female-targeted beauty economy that is pressurizing more and more faces to become socially “acceptable” while noisily touting the importance of individuality. The spatio-semantic recursion of *wanghong*, its move from people to places and from places back to people, further consolidates this sense of flattened diversity, as it enacts a process of distributed (de)personhood, in which the properties of human internet celebrity are first mapped onto urban space and then standardized, homogenized, and banalized still further by the same AI-driven processes that respond to the prompt “Show me a picture of a young Chinese woman” with the image of an Internet celebrity face. To put this another way, what is the symbolic qualitative difference between an anywhere-nowhere *wanghong* café and an Internet celebrity face stripped of its distinctive DNA markers? As the cityscape and the facescape—the where and who—start to blur under a system of AI governance mostly designed and implemented by men, it will get harder to locate female agency and perhaps harder still to identify what it looks like.

GenAI bots now populate internets everywhere with images and texts generated at ever-escalating levels of efficiency. This sinister reality has been dubbed “cannibalism,” a term that deliberately pathologizes the spread of AI tools that are fed and watered on an AI diet.³⁴ As AI continues to replicate itself—with social media accounts first created by AI, then maintained by a steady flow of AI-generated posts and images, and later kept in play by pseudoengagement from other AI-managed accounts—conspiracy theories have proliferated which postulate that AI-generated content is now overtaking the human in cyberspace, leaving us with a “dead internet.” As Renzella and Rozova put it, “there is strong evidence social media is being manipulated by these inflated bots to sway public opinion with disinformation” about everything from mass shootings in the United States to the war in Ukraine.³⁵ The baleful collaboration between face-editing, deepfakes, and GenAI in Chinese online space sheds worrisome new light on this notion of the internet in its death throes, and thus horribly vulnerable to untruths. First, it reveals the determinant—and highly ironic—role that human faces (though perhaps humanoid facialogy is becoming an apt term) have played in enabling this GenAI infiltration. Second,

it suggests that the deepest, most darkly constitutive politics at work here may be the politics of gender. Third, it intimates that when the internet starts to die, other human environments, with their life-sustaining truths, may follow in its wake. Yet perhaps most worrisome of all is the brazen transparency of all this. These flattening, deadening moves are not especially secretive or arcane; they are easy to see through and all too obvious. As such, they present us not just with endlessly duplicated faces, but also with the crude and perplexingly unchallenged technology of duplication itself.

Notes

1. See Baidu Baike's "Net red face" entry, <https://baike.baidu.com/item/网红脸/22207554>.
2. Yang Yidan 杨一丹, "'Renti zhadan' he 'wanghonglian': 'ququanqiuhua' yujingxia de shenti zhengzhi" "人体炸弹" 和 "网红脸:" "去全球化" 语境下的身体政治 ["Human body bombs" and the "Internet celebrity face": body politics in the context of "deglobalization"], *Wenxue yu wenhua* 文学与文化 [Literature and cultural studies] no. 2 (2017): 71.
3. Anett Dippner, *Social Media Celebrities and Neoliberal Body Politics in China*, KFG Working Paper Series 91 (Berlin: Free University of Berlin, 2018), 3.
4. Lisa Nan, "How AI Technology Is Reshaping Beauty Retail in China," *Campaign Asia*, 19 July 2023, <https://www.campaignasia.com/article/how-ai-technology-is-reshaping-beauty-retail-in-china/485394>.
5. Rouyu Li et al., "An Artificial Intelligence Tool Used for Patient Selection in Cosmetic Surgery," *Aesthetic Plastic Surgery* 49, no. 19 (2025): 5631–5638; Yi-Xin Sun et al., "GPT-4: The Future of Cosmetic Procedure Consultation?" *Aesthetic Surgery Journal* 43, no. 8 (2023): NP670–NP672; Rouyu Li et al. "Artificial Intelligence for Rhinoplasty Design in Asian Patients," *Aesthetic Plastic Surgery* 48, no. 8 (2024): 1557–1564; Yujin Kang, Junming Liu, and Haiyan Cui, "AI-Driven Assessment of Lip Volume Improvement Using Hyaluronic Acid Fillers: A Comprehensive Analysis," *Aesthetic Plastic Surgery* 49, no. 15 (2025): 4384–4391.
6. Liu Chunshan 刘春山 and Han Yang 韩阳, "Meitu xiuxiu zai tongyi dajia de shenmei? Meitu chuangshiren Wu Xinhong: 'Wanghonglian' yijing jinhua cheng 'qianren qianmian'" 美图秀秀在统一大家的审美? 美图创始人吴欣鸿: "网红脸"已经进化成"千人千面" [Is Meitu Xiuxiu unifying everyone's aesthetics? Meitu founder Wu Xinhong says: "'Internet celebrity face' has evolved into 'thousands of faces for thousands of different people'"], *Meiri jingji xinwen* 每日经济新闻 [Daily Economic News], 31 August 2021, <https://www.nbd.com.cn/articles/2021-08-31/1898311.html>.
7. On its company webpage, Meitu describes itself as having been an AI-driven company ever since its founding in 2008. See "Meitu, Inc.: History, Ownership, Mission, How It Works and Makes Money," Meitu, 20 February 2025, <https://www.meitu.com/hk/introduction>.
8. One notable difference is that a rule-based approach back in the early 2000s was unable to take advantage of more data.
9. Meitu's mobile app was launched in 2011, five years after Geoffrey Hinton, Simon Osindero, and Yee-Whye Teh published the paper that defined the field of deep learning. See Hinton, Osindero, and Teh, "A Fast Learning Algorithm for Deep Belief Nets," *Neural Computation* 18, no. 7 (2006): 1527–1554.
10. A slew of similar apps followed soon after, including VideoMaster (2013), AI Face Swap Photo Editor (2013), and Celebrity Face Swap (2014).
11. See Facebook's official launch video for Instagram face filters (16 May 2017), <https://www.facebook.com/instagram/videos/today-were-introducing-face-filters-an-easy-way-to-turn-an-ordinary-selfie-into-/1339813909437783/>.
12. Evan Jacoby, "I Paid \$30 to Create a Deepfake Porn of Myself," *Vice*, 9 December 2019, <https://www.vice.com/en/article/i-paid-dollar30-to-create-a-deepfake-porn-of-myself/>.
13. Although Reddit and X (formerly Twitter) have banned deepfaked porn, 4chan, 8chan, and other dedicated porn-sharing sites continue to help its spread. Richard Waters, "Twitter, Reddit Vow to Clamp Down on Fake Porn Videos," *Financial Times*, 7 February 2018, <https://www.ft.com/content/21b3fba0-0c52-11e8-8eb7-42f857ea9f09>.
14. See Justin Hendrix, "Transcript: TikTok CEO Testifies to Congress," *Tech Policy*, 25 March 2023, <https://www.techpolicy.press/transcript-tiktok-ceo-testifies-to-congress/>. On legal events, see "Mark Zuckerberg Opening

Statement Transcript: House Hearing on Misinformation,” *rev* (blog), 25 March 2021, <https://www.rev.com/transcripts/mark-zuckerberg-opening-statement-transcript-house-hearing-on-misinformation>.

15. We generated the outputs with the prompt “Generate an image of a young Chinese woman” using the models’ API. The ten outputs from each model are statistically independent.

16. Walter Benjamin, *One Way Street and Other Writings*, trans. J.A. Underwood (London: Penguin, 2009), 248.

17. Francis Galton, “Photographic Composites,” *Photographic News*, 17 April 1885, 243–5.

18. Chief among these artists-technicians was Chen Shilin, who revealed the previously clandestine details of his retouching process in an interview with Wang Shiyao 王诗瑶, “Biaojunzhao hou: fang Chen Shilin” 标准照后：访陈石林 [Behind the standard photograph: An interview with Chen Shilin], *Chinese Photographers (Zhongguo sheyingjia 中国摄影家)* 12 (2007): 64–67.

19. See Kathy Davis, *Reshaping the Female Body: The Dilemma of Cosmetic Surgery* (New York/London: Routledge, 1995), 56; and Susan Bordo, *Twilight Zones: The Hidden Life of Cultural Images from Plato to OJ* (Berkeley and Los Angeles: University of California Press, 1997), 37–44.

20. Kathryn Pauly Morgan, “Woman and the Knife: Cosmetic Surgery and the Colonization of Women’s Bodies,” *Hypatia* 6, no. 3 (1991): 36.

21. As Ana Sofia Elias and Rosalind Gill put it, “[I]t is clear that being invited to see how your face would look like after rhinoplasty or eyelid surgery while you are standing at the bus stop or waiting in line at the ATM radically changes the meanings of such interventions, rendering cosmetic surgery as more familiar, banal, and culturally intelligible as ‘normal.’” See Elias and Gill, “Beauty Surveillance: The Digital Self-monitoring Cultures of Neoliberalism,” *European Journal of Cultural Studies* 21, no. 1 (2017): 75.

22. This is a procedure commonly discussed on Chinese-language plastic surgery blogs and advertisement websites; Western plastic surgery advertisements, by contrast, are more likely to feature the opposite procedure (narrowing or removing part of the jawbone).

23. “Taobao and Tmall Launch AI Image-to-Video Function, Greatly Improving Merchant’s Video Production Efficiency!,” *Alibase*, 15 May 2025, <https://www.aibase.com/news/18093>.

24. Hadar Levy-Landesberg and Xuenan Cao, “Anchoring Voices: The News Anchor’s Voice in China from Television to AI,” *Media, Culture & Society* 47, no. 2 (2025): 229–251.

25. Yuanjie Xia, “AI Sign Language Anchor Serves at Olympics,” *Shenzhen Economic Daily*, 8 February 2022, https://www.eyeshenzhen.com/content/2022-02/08/content_24921512.htm.

26. Alibaba Alicloud provides a one-stop solution to help companies create AI influencers and avatars. See “Virtual Digital Person,” on the Alicloud website, accessed 19 November 2025, https://cn.aliyun.com/product/ai/avatar?from_alibabacloud=.

27. UNESCO, “Generative AI: UNESCO Study Reveals Alarming Evidence of Regressive Gender Stereotypes,” 7 March 2024, <https://www.unesco.org/en/articles/generative-ai-unesco-study-reveals-alarming-evidence-regressive-gender-stereotypes>.

28. See Chapter 4 of Janet Abbate, *Recoding Gender: Women’s Changing Participation in Computing* (Cambridge, MA: MIT Press, 2017) for testimonies from American women programmers in the 1960s. A very similar historical narrative can be found in Mar Hicks, *Programmed Inequality: How Britain Discarded Women Technologists and Lost Its Edge in Computing* (Cambridge, MA: MIT Press, 2018).

29. “Gender Gap in China’s Science Sector Narrowing, Report Suggest,” *Sixth Tone*, 4 November 2022, <https://www.sixthtone.com/news/1011577>.

30. Amy Y. Zhang, Asa Roast, and Carwyn Morris, “*Wanghong* Urbanism:

Towards a New Urban-Digital Spectacle,” *Mediapolis* 7, no. 4 (2022), <https://www.mediapolisjournal.com/2022/11/wanghong-urbanism/>.

31. Dino Ge Zhang, “The Aesthetics and Logic of *Wanghong* in Postdigital China,” in *Asian Celebrity Cultures in the Digital Age*, ed. Jian Xu, Glen Donnar, and Divya Garg (Hong Kong: Hong Kong University Press, 2025), 210.

32. Gertrude Stein, *Everybody’s Autobiography* (New York: Random House, 1937), 289.

33. Zhang, Roast, and Morris, “*Wanghong* Urbanism.”

34. One of the earliest critiques of generative AI using the term “cannibalism” is Melissa Heikkilä, “How AI-generated Text Is Poisoning the Internet,” *MIT Technology Review*, 20 December 2022, <https://www.technologyreview.com/2022/12/20/1065667/how-ai-generated-text-is-poisoning-the-internet/>. The term has since been popularized by researchers and mainstream news media.

35. Jake Renzella and Vlada Rozova, “The ‘Dead Internet Theory’ Makes Eerie Claims about an AI-Run Web: The Truth Is More Sinister,” *The Conversation*, 19 May 2024, <https://theconversation.com/the-dead-internet-theory-makes-eerie-claims-about-an-ai-run-web-the-truth-is-more-sinister-229609>.