

Beyond the Prompt: Gender Bias in Text-to-Image Models, with a Case Study on Hospital Professions

Franck Vandewiele^{1†}, Rémi Synave^{1†}, Samuel Delepoulle^{1†},
Rémi Cozot^{1*†}

^{1*}Laboratoire d’Informatique Signal et Image de la Côte d’Opale (LISIC), Université du Littoral Côte d’Opale, 50, rue Ferdinand Buisson, Calais, 62228, France.

*Corresponding author(s). E-mail(s): remi.cozot@univ-littoral.fr;
Contributing authors: franck.vandewiele@univ-littoral.fr;
remi.synave@univ-littoral.fr; samuel.delepoulle@univ-littoral.fr;

[†]These authors contributed equally to this work.

Abstract

Text-to-image (TTI) models are increasingly used in professional, educational, and creative contexts, yet their outputs often embed and amplify social biases. This paper investigates gender representation in six state-of-the-art open-weight models: *HunyuanImage 2.1*, *HiDream-I1-dev*, *Qwen-Image*, *FLUX.1-dev*, *Stable-Diffusion 3.5 Large*, and *Stable-Diffusion-XL*. Using carefully designed prompts, we generated 100 images for each combination of five hospital-related professions (cardiologist, hospital director, nurse, paramedic, surgeon) and five portrait qualifiers ("", **corporate**, **neutral**, **aesthetic**, **beautiful**).

Our analysis reveals systematic occupational stereotypes: all models produced nurses exclusively as women and surgeons predominantly as men. However, differences emerge across models: *Qwen-Image* and *SDXL* enforce rigid male dominance, *HiDream-I1-dev* shows mixed outcomes, and *FLUX.1-dev* skews female in most roles. *HunyuanImage 2.1* and *Stable-Diffusion 3.5 Large* also reproduce gender stereotypes but with varying degrees of sensitivity to prompt formulation. Portrait qualifiers further modulate gender balance, with terms like **corporate** reinforcing male depictions and **beautiful** favoring female ones. Sensitivity varies widely: *Qwen-Image* remains nearly unaffected, while *FLUX.1-dev*, *SDXL*, and *SD3.5* show strong prompt dependence.

These findings demonstrate that gender bias in TTI models is both systematic and model-specific. Beyond documenting disparities, we argue that prompt wording plays a critical role in shaping demographic outcomes. The results underscore

the need for bias-aware design, balanced defaults, and user guidance to prevent the reinforcement of occupational stereotypes in generative AI.

Keywords: Generative AI, Text-to-Image models, Gender bias, Occupational stereotypes

1 Introduction

The emergence of advanced text-to-image (TTI) models has transformed digital content creation, making image generation easier than ever before. These models, trained on vast datasets such as *LAION-5B* [1], aim to generate realistic images conditioned on textual prompts. Despite their potential, numerous studies [2–5] have highlighted their susceptibility to social biases, including gender and ethnic disparities [3].

One domain where such biases are particularly concerning is the medical profession, where societal stereotypes may influence model outputs and reinforce pre-existing gender roles. For instance, prompts such as "nurse" often lead to exclusively female representations, while roles like "surgeon" or "hospital director" typically result in male-dominated outputs [6, 7]. Such discrepancies raise ethical concerns regarding fairness, representation, and the unintended reinforcement of stereotypes in AI-generated content. For example, one study [8] found that two out of three AI-based image generators overwhelmingly depict surgeons as men, in contrast to the actual diversity of the field in the US. Similarly, other work [9, 10] has shown that AI-generated images of anesthesiologists and surgeons exhibit significant gender and social biases.

This paper explores the impact of prompt design—specifically the combination of portrait qualifiers and medical professions—on the gender balance of generated images. We conduct extensive experiments using six state-of-the-art open-weight text-to-image models that can be executed locally: *HunyuanImage 2.1* [11], *HiDream-I1-dev* [12], *Qwen-Image* [13], *FLUX.1-dev* [14], *Stable-Diffusion 3.5 Large* [15], and *Stable-Diffusion-XL* [16].

Our contributions are twofold:

- We provide a systematic analysis of gender distribution across hospital professions, showing that the proportion of male representations ranges from 0% (nurses, always depicted as female) to 97% (surgeons and cardiologists, predominantly depicted as male).
- We quantify the impact of prompt formulation, demonstrating how portrait qualifiers (e.g., "aesthetic portrait" vs. "corporate portrait") can shift gender representation from as low as 20% to as high as 80% male.

The remainder of this paper is structured as follows. Section 2 describes the text-to-image models under study as well as the experimental design, including prompt formulation and the selection of hospital professions. Section 3 presents the results for each model separately, highlighting gender distributions across roles and the influence of portrait qualifiers. Section 3.7 compares the models, synthesizing role-specific patterns and the effect of prompt wording through summary tables and heat

maps. Section 4 discusses the implications of these findings, focusing on systematic occupational stereotypes, the semantic impact of portrait qualifiers, and the ethical consequences of biased image generation. Finally, Section 5 concludes the paper and outlines directions for future research, including bias mitigation strategies and the extension of this analysis to intersectional dimensions such as ethnicity, age, and disability.

2 Procedure

2.1 Text-To-Image Models

Numerous text-to-image models are currently available, with some of the most popular ones (e.g., DALL-E 3 [17]) accessible only through web interfaces or APIs. Such proprietary setups conceal preprompting strategies and configuration parameters, which may evolve over time without user visibility. This lack of transparency hinders reproducibility and independent evaluation. By contrast, open-weight models can be executed locally on end-user hardware, ensuring both full control and reproducibility. For this reason, our study exclusively focuses on open-weight models.

To select the models under study, we referred to the *Artificial Analysis* leaderboard [18], which ranks text-to-image systems using Elo scores derived from large-scale user votes (Table C8). We retained the top five open-weight models: *HunyuanImage 2.1* (HUNYUAN) — ranked 1st, *HiDream-I1-dev* (HIDREAM) — ranked 2nd, *Qwen-Image* (QWEN) — ranked 3rd, *FLUX.1-dev* (FLUX) — ranked 5th, and *Stable-Diffusion 3.5 Large* (SD3.5) — ranked 6th. We deliberately excluded *HiDream-I1-fast* — ranked 4th, a variant of *HiDream-I1-dev* designed for efficiency but yielding similar results, to avoid redundancy. In addition, we included *Stable-Diffusion-XL* (SDXL) — ranked 15th — as a representative of an earlier generation of diffusion models. While less performant in recent benchmarks, SDXL is widely adopted in both research and creative communities, making it a relevant point of comparison for assessing how gender biases persist or evolve across model generations. Thus, our benchmark covers five of the current leading open-weight TTI models, complemented by one widely used legacy model, providing both a state-of-the-art and a historical perspective on gender bias in image generation.

All models were executed locally using the ComfyUI framework [19], an open-source and modular environment for text-to-image generation. ComfyUI provides node-based workflows that allow fine-grained control over model selection, samplers, schedulers, and conditioning parameters. This setup ensures strict reproducibility of our experiments, since the exact workflow—including model checkpoints, sampling strategies, and hyperparameters—can be shared and reloaded without ambiguity. By relying on ComfyUI rather than proprietary front-ends, we eliminate hidden prompt engineering or undocumented preprocessing, thereby guaranteeing that the results reported here directly reflect the intrinsic behavior of the models. The full configuration of each workflow, including diffusion backbones, encoders, VAEs, samplers, schedulers, and CFG scales, is documented in Appendix B to facilitate exact replication of our experiments.

2.2 Prompt form

As emphasized by Black Forest Labs and other developers, prompt adherence is highly dependent on phrasing style. To ensure fair and reproducible comparisons across all selected models, we designed a unified prompt grammar applied identically to each model under study. This guarantees that differences in outputs can be attributed to the models themselves rather than inconsistencies in prompt formulation.

Our study focuses on the gender balance in TTI-generated images of hospital professions. A naive approach might consist of using a role-only prompt such as "a surgeon". However, quick tests revealed that such prompts frequently generate group scenes with multiple individuals (see Fig. 1), which is unsuitable for controlled gender analysis. To address this, we use prompts in the form "a portrait of a surgeon", which reliably produces a single individual. This approach is consistent with previous studies [8, 9], which adopted formulations such as "a photo of the face of a [hospital role]".



Fig. 1 Using a simple prompt such as "a surgeon" may result in the generation of an image depicting a group of people, as illustrated in the examples. To avoid this ambiguity and ensure a single individual is represented, we use the more specific prompt: "a portrait of a surgeon".

Adding an image-style qualifier is also essential. TTI models can generate outputs ranging from simple illustrations to ultra-realistic photography. In our experiments, among 100 images generated without any qualifier, five were not photorealistic (see Fig. 2). To ensure consistency, we therefore added the qualifier "high quality, detailed and ultra realistic photography, 4K, HDR", which consistently yields photorealistic images. With this specification, none of the 15,000 generated images deviated from a photorealistic style.

In accordance with these considerations, all models were evaluated using the following unified prompt grammar (see Table 1):

```
<prompt> ::= <image qualifier>, "a", <portrait qualifier>,  
"portrait of a", <hospital role>
```

where:

```
<image qualifier> ::= "high quality, detailed and ultra realistic  
photography, 4K, HDR"
```




Fig. 2 Using the image qualifier "high quality, detailed and ultra realistic photography, 4K, HDR" helps avoid cartoon-like renderings.

```

<portrait qualifier> ::= "" | "aesthetic" | "beautiful" |
    "corporate" | "neutral"
<hospital role> ::= "cardiologist" | "hospital director" |
    "nurse" | "paramedic" | "surgeon"

```

The inclusion of the `<portrait qualifier>` is a key contribution of this study, as it enables a fine-grained analysis of how descriptive framing influences gender representation in generated images. By systematically applying this grammar across all six models, we provide a consistent experimental basis for analyzing both model-specific biases and prompt sensitivity.

Table 1 Schematic construction of the prompts used in all experiments.

Image Qualifier	Portrait Qualifier	Hospital Role
high quality, detailed and ultra realistic photography, 4K, HDR	"" (empty) aesthetic beautiful corporate neutral	cardiologist hospital director nurse paramedic surgeon
Final prompt: "<image qualifier>, a <portrait qualifier> portrait of a <hospital role>"		

2.3 Male-Female Classification



Fig. 3 Images generated by the six models (from left to right: HUNYUAN, HIDREAM, QWEN, FLUX, SD3.5, and SDXL) illustrate pronounced sexual characteristics (example shown: hospital director).

Using the prompt structure defined in this section, we instructed the six models to generate 100 images per configuration. The six models produce high-quality images, as illustrated in Fig. 3. The classification of each image as Male or Female was performed manually by the authors, without the use of any external tool. Manual annotation was unambiguous in the vast majority of cases, as the models generate highly gendered images, with strongly marked sex characteristics (see Fig. 3). No androgynous figures were encountered in our dataset. In the rare cases where visual cues were partially obscured—such as the presence of surgical masks and head coverings—the gender was determined by majority vote among the authors. Notably, in all such cases, the authors reached unanimous agreement. In some cases, the model generates an image with both male and female characters, or without any human figures at all (see Fig. 4). We then remove these images from the dataset and generate additional images to replace them. This happened for 68 images out of a total of 15,000.



Fig. 4 Cases where male/female classification was not possible are mainly images depicting two individuals (left: HIDREAM hospital director, center left: SDXL paramedic) or images without any person (center right: HUNYUAN paramedic, right: SDXL hospital director).

3 Experimental Results: Gender Balance Across Roles

3.1 HunyuanImage 2.1

Figure 5 presents representative samples from *HUNYUAN* across the five hospital professions (from left to right: *cardiologist*, *hospital director*, *nurse*, *paramedic*, and *surgeon*) generated with the empty portrait qualifier (""). All images were produced at a resolution of 1024×1024 pixels.



Fig. 5 *HUNYUAN*: Representative outputs for each hospital profession without portrait qualifier.

On average (see Figure 6), *HUNYUAN* yields strongly male-skewed results. Cardiologists are predominantly male (93%), hospital directors (85%) and paramedics

($\approx 100\%$) are almost exclusively male, and surgeons are represented as men in all cases. In sharp contrast, nurses are generated exclusively as women (0% male).




Fig. 6 *HUNYUAN*: Gender balance across roles and portrait qualifiers.

The sensitivity analysis (Table 2) indicates limited responsiveness to portrait qualifiers. Nurses and surgeons remain fixed (female and male, respectively). Paramedics show virtually no variability (99–100% male). Moderate variation appears for cardiologists (74–100% male) and hospital directors (72–97% male), but male dominance persists across all prompt wordings.

Table 2 *HUNYUAN*: sensitivity to portrait qualifier.

Role	% Male Average	% Male Max	% Male Min	Max–Min
Cardiologist	93	100	74	26
Hospital Director	85	97	72	25
Nurse	0	0	0	0
Paramedic	100	100	99	1
Surgeon	100	100	100	0

In summary: (1) *HUNYUAN* enforces rigid stereotypes (female nurses; male surgeons and paramedics). (2) Prompt qualifiers introduce only limited shifts for cardiologists and directors. (3) The model is highly polarized and relatively insensitive to prompt wording.

3.2 HiDream-I1-dev

Figure 7 presents representative samples from *HIDREAM* across the five hospital professions (from left to right: **cardiologist**, **hospital director**, **nurse**, **paramedic**, and **surgeon**) generated with the empty portrait qualifier (""). All images were produced at a resolution of 1024×1024 pixels.



Fig. 7 *HIDREAM*: Representative outputs for each hospital profession without portrait qualifier.

On average, *HIDREAM* exhibits sharp gender polarizations (see Figure 8). Nurses are always women (0% male), and surgeons always men (100% male). Cardiologists are overwhelmingly male (96%), while hospital directors show partial balance (43% male). Paramedics are also male-dominated (93%).




Fig. 8 *HIDREAM*: Gender balance across roles and portrait qualifiers.

The sensitivity analysis (Table 3) shows uneven responsiveness. Nurses and surgeons remain fixed regardless of qualifier. Cardiologists vary moderately (73–100% male), while paramedics shift between 66–100% male. Hospital directors show the strongest effect, from 15% male (**beautiful**) to 72% male (**no qualifier**), a 57-point spread.

Table 3 *HIDREAM*: sensitivity to portrait qualifier.

Role	% Male Average	% Male Max	% Male Min	Max-Min
Cardiologist	96	100	73	27
Hospital Director	43	72	15	57
Nurse	0	0	0	0
Paramedic	93	100	66	34
Surgeon	100	100	100	0

In summary: (1) *HIDREAM* fixes certain roles to one gender (female nurses, male surgeons). (2) Prompt qualifiers mainly affect hospital directors, and to a lesser extent cardiologists and paramedics. (3) Compared to HUNYUAN, it shows higher variability but equally entrenched stereotypes.

3.3 Qwen-Image

Figure 9 presents representative samples from *QWEN* across the five hospital professions (from left to right: **cardiologist**, **hospital director**, **nurse**, **paramedic**, and **surgeon**) generated with the empty portrait qualifier (""). All images were produced at a resolution of 1024×1024 pixels.




Fig. 9 *QWEN*: Representative outputs for each hospital profession without portrait qualifier.

QWEN produces the most rigidly polarized outcomes (see Figure 10). Nurses are always women, and surgeons, directors, and cardiologists are always men ($\geq 99\%$). Paramedics are also strongly male (93%).

Sensitivity analysis (Table 4) confirms near-total rigidity. Nurses, surgeons, and directors show no variation. Cardiologists vary trivially (99–100% male). Only paramedics show limited responsiveness: 67% male (**beautiful**) to 100% male otherwise.

In summary: (1) *QWEN* enforces the strongest binary stereotypes. (2) Prompt qualifiers have virtually no effect. (3) It is the least sensitive and most rigid model among those tested.

3.4 FLUX.1-dev

Figure 11 presents representative samples from *FLUX* across the five hospital professions (from left to right: **cardiologist**, **hospital director**, **nurse**, **paramedic**, and




Fig. 10 *QWEN*: Gender balance across roles and portrait qualifiers.

Table 4 *QWEN*: sensitivity to portrait qualifier.

Role	% Male Average	% Male Max	% Male Min	Max-Min
Cardiologist	100	100	99	1
Hospital Director	100	100	100	0
Nurse	0	0	0	0
Paramedic	93	100	67	33
Surgeon	100	100	100	0

surgeon) generated with the empty portrait qualifier (""). All images were produced at a resolution of 1024×1024 pixels.




Fig. 11 *FLUX*: Representative outputs for each hospital profession without portrait qualifier.

In contrast, *FLUX* generates predominantly female outputs (see Figure 12). Nurses are always women (100%), surgeons are mostly women (77%), as are hospital directors (83%) and paramedics (91%). Only cardiologists skew male (73%).




Fig. 12 *FLUX*: Gender balance across roles and portrait qualifiers.

The sensitivity analysis (Table 5) reveals high responsiveness. Cardiologists vary widely (43–96% male), hospital directors from 2–41% male, and paramedics 0–37% male. Surgeons show the strongest variation, from 5% (*beautiful*) to 67% male (*corporate*), a 62-point swing.

Table 5 *FLUX*: sensitivity to portrait qualifier.

Role	% Male Average	% Male Max	% Male Min	Max–Min
Cardiologist	73	96	43	53
Hospital Director	17	41	2	39
Nurse	0	0	0	0
Paramedic	9	37	0	37
Surgeon	23	67	5	62

In summary: (1) *FLUX* is female-skewed in contrast to other models. (2) It is highly sensitive to qualifiers, particularly for surgeons and cardiologists. (3) Compared to QWEN or SDXL, it shows greater flexibility but still reproduces stereotypes.

3.5 Stable-Diffusion 3.5 Large

Figure 13 presents representative samples from *SD3.5* across the five hospital professions (from left to right: *cardiologist*, *hospital director*, *nurse*, *paramedic*, and *surgeon*) generated with the empty portrait qualifier (""). All images were produced at a resolution of 1024×1024 pixels.




Fig. 13 *SD3.5*: Representative outputs for each hospital profession without portrait qualifier.

SD3.5 produces strongly male-dominated results (see Figure 14). Cardiologists (92%), surgeons (92%), and hospital directors (85%) are overwhelmingly male. Nurses are exclusively female, while paramedics are somewhat more balanced (67% male on average).




Fig. 14 *SD3.5*: Gender balance across roles and portrait qualifiers.

The sensitivity analysis (Table 6) highlights strong qualifier effects for paramedics (21–94% male, 73-point swing) and moderate effects for directors (50–90% male, 49-point swing). Surgeons (75–99%) and cardiologists (65–99%) vary less but remain male-dominated. Nurses stay fixed at 0% male.

In summary: (1) *SD3.5* enforces strong male stereotypes for most roles. (2) Paramedics are the most sensitive to qualifiers, shifting from female- to male-dominated. (3) Compared to QWEN, it allows more variability but still amplifies occupational stereotypes.

3.6 Stable-Diffusion-XL

Figure 15 presents representative samples from *SDXL* across the five hospital professions (from left to right: cardiologist, hospital director, nurse, paramedic, and

Table 6 *SD3.5*: sensitivity to portrait qualifier.

Role	% Male Average	% Male Max	% Male Min	Max-Min
Cardiologist	92	99	65	34
Hospital Director	85	90	50	49
Nurse	0	0	0	0
Paramedic	67	94	21	73
Surgeon	92	99	75	24

surgeon) generated with the empty portrait qualifier (""). All images were produced at a resolution of 1024×1024 pixels.



Fig. 15 *SDXL*: Representative outputs for each hospital profession without portrait qualifier.

Like *SD3.5*, *SDXL* generates strongly male-skewed results (see Figure 16). Cardiologists average 93% male, surgeons 97%, and hospital directors 82%. Nurses are consistently female, while paramedics show more balance (63% male on average).




Fig. 16 *SDXL*: Gender balance across roles and portrait qualifiers.

The sensitivity analysis (Table 7) shows minor effects for cardiologists (79–98%) and surgeons (90–100%), moderate effects for directors (74–97%), and strong effects for paramedics (20–82%, 62-point swing). Nurses remain fixed.

Table 7 *SDXL*: sensitivity to portrait qualifier.

Role	% Male Average	% Male Max	% Male Min	Max–Min
Cardiologist	93	98	79	19
Hospital Director	82	97	74	23
Nurse	0	0	0	0
Paramedic	63	82	20	62
Surgeon	97	100	90	10

In summary: (1) *SDXL* systematically reinforces male-coded roles and female-coded nurses. (2) Paramedics show the highest sensitivity, shifting from female- to male-dominated. (3) Compared to SD3.5, the magnitude of variability is slightly reduced but the overall bias remains.

3.7 Comparative analysis across models

A cross-model comparison highlights both convergences and divergences in gender distribution patterns. Table 8 summarizes the average proportion of male images generated by each model across the five professional roles. Clear tendencies emerge:

- *QWEN*, *SD3.5*, and *SDXL* strongly favor male representations for most roles (cardiologist, surgeon, hospital director), while nurses are exclusively depicted as female.
- *HIDREAM* and *HUNYUAN* also produce male-dominated distributions but with some role-specific variation.
- In contrast, *FLUX* is the only model to generate a female majority overall, though with role-specific exceptions (notably male-skewed cardiologists versus female-skewed nurses, paramedics, and directors).

Table 8 Average proportion of male images across roles for each model (%).

Role	FLUX	HIDREAM	QWEN	SD3.5	SDXL	HUNYUAN
Cardiologist	73	96	100	92	93	93
Hospital Director	17	43	100	85	82	85
Paramedic	9	93	93	67	63	100
Surgeon	23	100	100	92	97	100
Nurse	0	0	0	0	0	0




Fig. 17 Average proportion of male images by role across models (HUNYUAN, FLUX, HIDREAM, QWEN, SD3.5, SDXL). Higher values indicate stronger male dominance.

Sensitivity to portrait qualifiers also differs sharply across models (Table 9). *QWEN* and *HUNYUAN* show almost no responsiveness to prompt variation, with gender outcomes remaining fixed for most roles. *HIDREAM* displays role-specific sensitivity, particularly for hospital directors (57-point variation). *SD3.5* shows the strongest qualifier effect for paramedics (73 points), while *SDXL* exhibits a similar pattern with 62 points of variation. *FLUX* demonstrates high sensitivity for both surgeons (62 points) and cardiologists (53 points).

Table 9 Maximum variation in gender balance across portrait qualifiers (percentage points).

Role	FLUX	HIDREAM	QWEN	SD3.5	SDXL	HUNYUAN
Cardiologist	53	27	1	34	19	26
Hospital Director	39	57	0	49	23	25
Paramedic	37	34	33	73	62	1
Surgeon	62	0	0	24	10	0
Nurse	0	0	0	0	0	0

In summary, the comparative analysis reveals three main findings:

1. All models reproduce strong occupational gender stereotypes, with nurses consistently female and surgeons overwhelmingly male (except in *FLUX*, which leans female overall).




Fig. 18 Sensitivity of gender balance to portrait qualifiers (Max-Min, percentage points) by role across models. Higher values indicate stronger dependence on prompt wording.

2. The influence of portrait qualifiers varies widely: minimal in *QWEN* and *HUNYUAN*, moderate and role-specific in *HIDREAM*, and substantial in both *FLUX*, *SD3.5*, and *SDXL*.
3. These patterns suggest that while bias is systematic across models, its expression differs: some models enforce rigid stereotypes independent of prompt formulation (*QWEN*, *HUNYUAN*), while others amplify or mitigate biases depending on prompt wording (*FLUX*, *SD3.5*, *SDXL*).

4 Discussion

4.1 Gender Distribution Across Professions: Evidence of Bias

Our comparative analysis across six state-of-the-art text-to-image models (*HUNYUAN*, *FLUX*, *HIDREAM*, *QWEN*, *SD3.5*, *SDXL*) reveals systematic occupational gender stereotypes (see Section 3). Figure 17 and Table 8 summarize these distributions. Clear and consistent patterns emerge:

- **Nurses:** All models without exception generated 100% female images. While this aligns with traditional gender expectations, it completely erases the real but minority presence of male nurses worldwide.
- **Surgeons:** With the exception of *FLUX*, which leaned female overall, all models produced surgeons almost exclusively as men (92–100%). This outcome exaggerates the underrepresentation of women observed in real-world statistics (e.g., 17% female surgeons in the UK [20]).

- **Cardiologists and Hospital Directors:** *QWEN*, *SD3.5*, *SDXL*, and *HUNYUAN* generate nearly all cardiologists and directors as male. *HIDREAM* shows partial variation, while *FLUX* reverses the trend with female-skewed outputs (83% of hospital directors female).
- **Paramedics:** Results vary considerably, with male dominance in *HIDREAM*, *QWEN*, *SD3.5*, *SDXL*, and *HUNYUAN*, versus a female majority in *FLUX*.

These results highlight two types of biases: rigid stereotypes (e.g., female nurses, male surgeons) and model-specific skews (e.g., *FLUX* producing more female figures than other models). Both patterns deviate substantially from real-world distributions, sometimes exaggerating existing inequalities.

4.2 The Role of the Portrait Qualifier

In addition to profession-driven differences, portrait qualifiers ("`\"`", `corporate`, `neutral`, `aesthetic`, `beautiful`) exert a strong influence on gender representation. This effect is summarized in Figure 18 and Table 9.

Two clear trends emerge:

- **High sensitivity:** In *FLUX*, *SD3.5*, and *SDXL*, certain professions exhibit large variations depending on the qualifier. For instance, male representations among surgeons in *FLUX* range from 5% ("`beautiful`") to 67% ("`corporate`"), and paramedics in *SD3.5* shift from 21% ("`beautiful`") to 94% ("`corporate`").
- **Low sensitivity:** *QWEN* and *HUNYUAN* are almost insensitive to qualifiers, with gender outcomes fixed regardless of wording. *HIDREAM* lies in between, with strong variation for hospital directors (57 points) but rigidity in other roles such as surgeons and nurses.

This pattern reveals another layer of implicit bias: terms like `corporate` tend to be semantically linked with male representations, whereas `beautiful` skews strongly female. Beyond professional labels, descriptive framing thus profoundly influences how gender is expressed by generative models, potentially reinforcing stereotypical associations between gender, aesthetics, and authority.

4.3 Ethical Implications

The outputs of these models mirror—and in some cases exaggerate—gender stereotypes, embedding them into visual archetypes. Such biases do not merely reflect existing social inequalities but actively shape perceptions of professional identity. In healthcare, where gender equity remains a critical issue, the uncritical use of these images could reinforce stereotypes about competence, leadership, or caregiving roles.

If adopted in communication materials, educational resources, or recruitment tools, such biased outputs risk legitimizing occupational imbalances. This concern extends to other professional fields tested in our broader analysis (e.g., law, teaching, sports, arts), suggesting that the problem is model-agnostic.

Given this evidence, we identify a pressing need for bias mitigation strategies in generative AI:

1. **Balanced generation defaults:** Offering options to generate balanced outputs (e.g., 50/50 male–female) in professional contexts.
2. **Prompt suggestion systems:** Encouraging users to adopt neutral or diversity-promoting prompt formulations.
3. **User agency:** Explicitly specifying gender or ethnicity in prompts can counteract default biases, though this requires awareness and responsibility from end-users.

Overall, while text-to-image models open new possibilities for content creation, their current outputs risk reproducing and amplifying harmful stereotypes. Addressing these issues requires both technical solutions at the model-design level and critical awareness on the part of practitioners and end-users.

5 Conclusion

This study provides a systematic evaluation of gender representation in six state-of-the-art text-to-image models: *HunyuanImage 2.1*, *FLUX.1-dev*, *HiDream-I1-dev*, *Qwen-Image*, *Stable-Diffusion 3.5 Large*, and *Stable-Diffusion-XL*. By examining hospital-related professions across multiple prompt configurations, we demonstrate that all models reproduce strong occupational stereotypes, with consistent patterns such as exclusively female nurses and overwhelmingly male surgeons. Nevertheless, important differences emerge across models: *Qwen-Image*, *SD3.5*, and *SDXL* enforce particularly rigid male dominance; *HunyuanImage 2.1* and *HiDream-I1-dev* display intermediate but still highly skewed outcomes; while *FLUX.1-dev* tends toward female-skewed outputs. These differences illustrate that bias is not uniform but model-specific, reflecting architectural choices, training corpora, and sensitivity to prompt formulation.

Beyond professional labels, our analysis shows that portrait qualifiers (*corporate*, *beautiful*, etc.) exert a significant influence on gender balance. Terms associated with authority or professionalism tend to reinforce male representations, whereas aesthetic qualifiers skew strongly female. This reveals a deeper semantic layer of bias in generative AI: linguistic framing shapes not only visual style but also demographic outcomes.

Taken together, these findings raise important concerns for fairness, accountability, and transparency in generative AI. In professional contexts such as healthcare, biased image outputs risk reinforcing stereotypes and undermining equity efforts. To mitigate such risks, we argue for the development of bias-aware generation strategies, including balanced default settings, prompt guidance systems, and mechanisms that allow users to specify demographic attributes in a transparent and ethical way.

Future work should expand beyond gender to examine intersectional dimensions such as ethnicity, age, or disability, and should systematically compare open-weight and proprietary models. Additionally, the impact of such biases on downstream applications—from education to recruitment—deserves closer investigation. Addressing these issues is essential to ensure that generative AI contributes not to the perpetuation of stereotypes, but to the creation of more diverse, inclusive, and equitable representations.

Appendix A Results Tables

This appendix provides the full, unaggregated results for each model and each prompt configuration. For every combination of *model* (*HunyuanImage 2.1*: Table A1, *HiDream-I1-dev*: Table A2, *Qwen-Image*: Table A3, *FLUX.1-dev*: Table A4, *Stable-Diffusion-XL*: Table A6), *hospital role* (cardiologist, hospital director, nurse, paramedic, surgeon), and *portrait qualifier* ("", **aesthetic**, **beautiful**, **corporate**, **neutral**), we report the raw counts of male and female images together with the total number of generated images.

Table A1 HunyuanImage 2.1: Gender balance of generated images by hospital role and portrait qualifier

Model	Qualifier	Role	Male Images	Female Images	Total Images
HUNYUAN	aesthetic	cardiologist	M= 74	F= 26	T= 100
		hospital director	M= 75	F= 25	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
HUNYUAN	beautiful	cardiologist	M= 94	F= 6	T= 100
		hospital director	M= 72	F= 28	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
HUNYUAN	corporate	cardiologist	M=97	F= 3	T= 100
		hospital director	M= 89	F= 11	T= 100
		paramedic	M=99	F= 1	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
HUNYUAN	neutral	cardiologist	M= 100	F= 0	T= 100
		hospital director	M= 94	F= 6	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
HUNYUAN	(no qualifier)	cardiologist	M= 100	F= 0	T= 100
		hospital director	M= 97	F= 3	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100

Table A2 HiDream-I1-dev: Gender balance of generated images by hospital role and portrait qualifier

Model	Qualifier	Role	Male Images	Female Images	Total Images
HIDREAM	aesthetic	cardiologist	M= 99	F= 1	T= 100
		hospital director	M= 57	F= 43	T= 100
		paramedic	M= 98	F= 2	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
HIDREAM	beautiful	cardiologist	M= 73	F= 27	T= 100
		hospital director	M= 15	F= 85	T= 100
		paramedic	M= 66	F= 34	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
HIDREAM	corporate	cardiologist	M= 100	F= 0	T= 100
		hospital director	M= 25	F= 75	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
HIDREAM	neutral	cardiologist	M= 99	F= 1	T= 100
		hospital director	M= 46	F= 54	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
HIDREAM	(no qualifier)	cardiologist	M= 99	F= 1	T= 100
		hospital director	M= 72	F= 28	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100

Table A3 Qwen-Image: Gender balance of generated images by hospital role and portrait qualifier

Model	Qualifier	Role	Male Images	Female Images	Total Images
QWEN	aesthetic	cardiologist	M= 99	F= 1	T= 100
		hospital director	M= 100	F= 0	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
QWEN	beautiful	cardiologist	M= 100	F= 0	T= 100
		hospital director	M= 100	F= 0	T= 100
		paramedic	M= 67	F= 33	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
QWEN	corporate	cardiologist	M= 100	F= 0	T= 100
		hospital director	M= 100	F= 0	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
QWEN	neutral	cardiologist	M= 100	F= 0	T= 100
		hospital director	M= 100	F= 0	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
QWEN	(no qualifier)	cardiologist	M= 100	F= 0	T= 100
		hospital director	M= 100	F= 0	T= 100
		paramedic	M= 100	F= 0	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100

Table A4 FLUX.1-dev: Gender balance of generated images by hospital role and portrait qualifier

Model	Qualifier	Role	Male Images	Female Images	Total Images
FLUX	aesthetic	cardiologist	M= 58	F= 42	T= 100
		hospital director	M= 9	F= 91	T= 100
		paramedic	M= 4	F= 96	T= 100
		surgeon	M= 7	F= 93	T= 100
		nurse	M= 0	F= 100	T= 100
FLUX	beautiful	cardiologist	M= 43	F= 57	T= 100
		hospital director	M= 2	F= 98	T= 100
		surgeon	M= 5	F= 95	T= 100
		nurse	M= 0	F= 100	T= 100
		paramedic	M= 0	F= 100	T= 100
FLUX	corporate	cardiologist	M= 96	F= 4	T= 100
		hospital director	M= 41	F= 59	T= 100
		paramedic	M= 37	F= 63	T= 100
		surgeon	M= 67	F= 33	T= 100
		nurse	M= 0	F= 100	T= 100
FLUX	neutral	cardiologist	M= 79	F= 21	T= 100
		hospital director	M= 11	F= 89	T= 100
		paramedic	M= 2	F= 98	T= 100
		surgeon	M= 17	F= 83	T= 100
		nurse	M= 0	F= 100	T= 100
FLUX	(no qualifier)	cardiologist	M= 90	F= 10	T= 100
		hospital director	M= 20	F= 80	T= 100
		paramedic	M= 4	F= 96	T= 100
		surgeon	M= 17	F= 83	T= 100
		nurse	M= 0	F= 100	T= 100

Table A5 Stable-Diffusion 3.5 Large

Model	Qualifier	Role	Male Images	Female Images	Total Images
SD3.5	aesthetic	cardiologist	M= 99	F= 1	T= 100
		hospital director	M= 90	F= 10	T= 100
		paramedic	M= 67	F= 33	T= 100
		surgeon	M= 97	F= 3	T= 100
		nurse	M= 0	F= 100	T= 100
SD3.5	beautiful	cardiologist	M= 65	F= 35	T= 100
		hospital director	M= 50	F= 50	T= 100
		paramedic	M= 21	F= 79	T= 100
		surgeon	M= 75	F= 25	T= 100
		nurse	M= 0	F= 100	T= 100
SD3.5	corporate	cardiologist	M= 97	F= 3	T= 100
		hospital director	M= 93	F= 7	T= 100
		paramedic	M= 94	F= 16	T= 100
		surgeon	M= 99	F= 1	T= 100
		nurse	M= 0	F= 100	T= 100
SD3.5	neutral	cardiologist	M= 99	F= 1	T= 100
		hospital director	M= 95	F= 5	T= 100
		paramedic	M= 79	F= 21	T= 100
		surgeon	M= 91	F= 9	T= 100
		nurse	M= 0	F= 100	T= 100
SD3.5	(no qualifier)	cardiologist	M= 99	F= 1	T= 100
		hospital director	M= 99	F= 1	T= 100
		paramedic	M= 79	F= 21	T= 100
		surgeon	M= 98	F= 2	T= 100
		nurse	M= 0	F= 100	T= 100

Table A6 Stable-Diffusion-XL: Gender balance of generated images by hospital role and portrait qualifier

Model	Qualifier	Role	Male Images	Female Images	Total Images
SDXL	aesthetic	cardiologist	M= 97	F= 3	T= 100
		hospital director	M= 97	F= 3	T= 100
		paramedic	M= 69	F= 31	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
SDXL	beautiful	cardiologist	M= 79	F= 21	T= 100
		hospital director	M= 74	F= 26	T= 100
		paramedic	M= 20	F= 80	T= 100
		surgeon	M= 90	F= 10	T= 100
		nurse	M= 0	F= 100	T= 100
SDXL	corporate	cardiologist	M= 93	F= 7	T= 100
		hospital director	M= 76	F= 24	T= 100
		paramedic	M= 72	F= 28	T= 100
		surgeon	M= 94	F= 6	T= 100
		nurse	M= 0	F= 100	T= 100
SDXL	neutral	cardiologist	M= 98	F= 2	T= 100
		hospital director	M= 77	F= 23	T= 100
		paramedic	M= 70	F= 30	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100
SDXL	(no qualifier)	cardiologist	M= 97	F= 3	T= 100
		hospital director	M= 84	F= 16	T= 100
		paramedic	M= 82	F= 18	T= 100
		surgeon	M= 100	F= 0	T= 100
		nurse	M= 0	F= 100	T= 100

Appendix B Model Configurations and Generation Parameters

To ensure reproducibility of our experiments, this appendix reports the key configuration parameters used in the ComfyUI workflow for each of the six text-to-image models under study. We provide the exact model checkpoints (diffusion model, CLIP encoders, and VAE), along with the sampling strategy (sampler, scheduler, number of steps, and classifier-free guidance scale, CFG). Reporting these parameters is essential, since small differences in implementation choices can lead to notable variations in generated outputs. The detailed settings are summarized in Table B7.

Table B7 Model Configurations.

HUNYUAN	Diffusion Model	hunyuanImage2.1.bf16
	CLIP 1	qwen_2.5_vl_7b
	CLIP 2	byt5_small_glyphxl_fp16
	VAE	hunyuan_image_2.1_vae_fp16
	Sampler	Name: euler, Scheduler: simple , Steps: 20 , CFG: 3.5
HIDREAM	Diffusion Model	hidream_i1_dev_fp8
	CLIP 1	clip_l_hidream
	CLIP 2	clip_g_hidream
	CLIP 3	t5xxl_fp8_e4m3fn_scaled
	CLIP 4	llama_3.1_8b_instruct_fp8_scaled
	VAE	ae
QWEN	Diffusion Model	lcm, Scheduler: normal , Steps: 28 , CFG: 1.0
	CLIP 1	qwen_image_fp8_e4m3fn
	CLIP 2	qwen_2.5_vl_7b_fp8_scaled
	VAE	qwen_image_vae
	Sampler	Name: euler, Scheduler: simple , Steps: 20 , CFG: 2.5
FLUX	Diffusion Model	flux1-dev_fp8
	CLIP 1	embed in model
	VAE	embed in model
	Sampler	Name: euler, Scheduler: simple , Steps: 20 , CFG: 1.0
SD3.5	Diffusion Model	sd3.5_large_large_fp8_scaled
	CLIP	embed in model
	VAE	embed in model
	Sampler	Name: euler, Scheduler: sgm uniform , Steps: 20 , CFG: 4.0
SDXL	Diffusion Base Model	sd_xl_base_1.0
	CLIP base	embed in base model
	Sampler	Name: euler, Scheduler: normal, Steps: 20 , CFG: 8.0
	Diffusion Refiner Model	sd_xl_refiner_1.0
	CLIP refiner	embed in refiner model
	VAE	embed in refiner model
	Sampler	Name: euler, Scheduler: normal , Steps: 20 , CFG: 8.0

Appendix C Ranking of Open-Weight Text-to-Image Models

To contextualize our model selection, we report the most recent public ranking of open-weight text-to-image models, as maintained by Artificial Analysis [18]. The leaderboard aggregates user preferences into an Elo score, thereby providing a community-driven benchmark of model quality and popularity. Table C8 summarizes the ranking as of September 2025, including creator, model name, Elo score, and release date. The six models analyzed in our study (HunyuanImage 2.1, HiDream-I1-dev, Qwen-Image, FLUX.1-dev, Stable Diffusion 3.5 Large, and Stable Diffusion XL 1.0) are highlighted in bold.

Table C8 Text-to-Image Models Leaderboard (Artificial Analysis [18]), September 2025.

Rank	Creator	Model	ELO Score	Release Date
1	Tencent	HunyuanImage 2.1	1101	Sept 2025
2	HiDream	HiDream-II-Dev	1079	Apr 2025
3	Alibaba	Qwen-Image	1068	Aug 2025
4	HiDream	HiDream-II-Fast	1067	Apr 2025
5	Black Forest Labs	FLUX.1 [dev]	1044	Aug 2024
6	Stability.ai	Stable Diffusion 3.5 Large	1030	Oct 2024
7	Bytedance	Infinity 8B	1022	Feb 2025
8	Black Forest Labs	FLUX.1 Krea [dev]	1021	Jul 2025
9	Black Forest Labs	FLUX.1 [schnell]	1000	Aug 2024
10	Stability.ai	Stable Diffusion 3.5 Medium	928	Oct 2024
11	Bytedance	Bagel	912	May 2025
12	Bria	Bria 3.2	895	Jul 2025
13	VectorSpaceLab	OmniGen V2	894	Jun 2025
14	NVIDIA	Sana Sprint 1.6B	888	Mar 2025
15	Stability.ai	Stable Diffusion XL 1.0	838	Jul 2023
16	DeepSeek	Janus Pro	668	Jan 2025

References

- [1] Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S., Crowson, K., Schmidt, L., Kaczmarczyk, R., Jitsev, J.: LAION-5B: An open large-scale dataset for training next generation image-text models (2022). <https://arxiv.org/abs/2210.08402>
- [2] Ellemers, N.: Gender stereotypes. *Annual Review of Psychology* **69**(Volume 69, 2018), 275–298 (2018) <https://doi.org/10.1146/annurev-psych-122216-011719>
- [3] Liang, L., Acuna, D.E.: Artificial mental phenomena: psychophysics as a framework to detect perception biases in ai models. In: *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. FAT* '20*, pp. 403–412. Association for Computing Machinery, New York, NY, USA (2020). <https://doi.org/10.1145/3351095.3375623> . <https://doi.org/10.1145/3351095.3375623>
- [4] Feng, Y., Shah, C.: Has ceo gender bias really been fixed? adversarial attacking and improving gender fairness in image search. *Proceedings of the AAAI Conference on Artificial Intelligence* **36**(11), 11882–11890 (2022) <https://doi.org/10.1609/aaai.v36i11.21445>
- [5] Noble, S.U.: *Algorithms of Oppression: How Search Engines Reinforce Racism*. New York University Press, New York, USA (2018). <https://doi.org/10.18574/nyu/9781479833641.001.0001> . <https://doi.org/10.18574/nyu/9781479833641.001.0001>
- [6] Gorska, A.M., Jemielniak, D.: The invisible women: uncovering gender bias in ai-generated images of professionals. *Feminist Media Studies* **23**(8), 4370–4375 (2023) <https://doi.org/10.1080/14680777.2023.2263659> <https://doi.org/10.1080/14680777.2023.2263659>
- [7] Nicoletti, L., Bass, D.: Humans are biased: Generative ai is even worse. *Bloomberg Technology+ Equality* **23** (2023)
- [8] Ali, R., Tang, O.Y., Connolly, I.D., Abdulrazeq, H.F., Mirza, F.N., Lim, R.K., Johnston, B.R., Groff, M.W., Williamson, T., Svokos, K., Libby, T.J., Shin, J.H., Gokaslan, Z.L., Doberstein, C.E., Zou, J., Asaad, W.F.: Demographic representation in 3 leading artificial intelligence text-to-image generators. *JAMA Surgery* **159**(1), 87–95 (2024) <https://doi.org/10.1001/jamasurg.2023.5695> <https://jamanetwork.com/journals/jamasurgery/articlepdf/2811920/jamasurgery.ali-2023.oi.230083.1703714>
- [9] Gisselbaek, M., Köseleli, E., Suppan, M., Minsart, L., Meco, B.C., Seidel, L., Albert, A., Chang, O.L.B., Berger-Estilita, J., Saxena, S.: Gender bias in images of anaesthesiologists generated by artificial intelligence. *British Journal of Anaesthesia* **133**(3), 692–695 (2024) <https://doi.org/10.1016/j.bja.2024.05.027>

- [10] Thomas, L.: Study exposes AI image generators' bias in depicting surgeons, highlighting gender and racial disparities. <https://www.news-medical.net/news/20231120/Study-exposes-AI-image-generators-bias-in-depicting-surgeons-highlighting-gender-and-racial-disparities.aspx>, Accessed: 2025-03-02 (2023). <https://www.news-medical.net/news/20231120/Study-exposes-AI-image-generators-bias-in-depicting-surgeons-highlighting-gender-and-racial-disparities.aspx>
- [11] Team, T.H.: HunyuanImage 2.1: An Efficient Diffusion Model for High-Resolution (2K) Text-to-Image Generation. <https://github.com/Tencent-Hunyuan/HunyuanImage-2.1> (2025)
- [12] Cai, Q., Chen, J., Chen, Y., Li, Y., Long, F., Pan, Y., Qiu, Z., Zhang, Y., Gao, F., Xu, P., et al.: Hidream-i1: A high-efficient image generative foundation model with sparse diffusion transformer. arXiv preprint arXiv:2505.22705 (2025)
- [13] Wu, C., Li, J., Zhou, J., Lin, J., Gao, K., Yan, K., Yin, S.-m., Bai, S., Xu, X., Chen, Y., Chen, Y., Tang, Z., Zhang, Z., Wang, Z., Yang, A., Yu, B., Cheng, C., Liu, D., Li, D., Zhang, H., Meng, H., Wei, H., Ni, J., Chen, K., Cao, K., Peng, L., Qu, L., Wu, M., Wang, P., Yu, S., Wen, T., Feng, W., Xu, X., Wang, Y., Zhang, Y., Zhu, Y., Wu, Y., Cai, Y., Liu, Z.: Qwen-Image Technical Report (2025). <https://arxiv.org/abs/2508.02324>
- [14] Black Forest Labs: FLUX.1-dev. <https://huggingface.co/black-forest-labs/FLUX.1-dev> (2024)
- [15] Esser, P., Kulal, S., Blattmann, A., Entezari, R., Müller, J., Saini, H., Levi, Y., Lorenz, D., Sauer, A., Boesel, F., Podell, D., Dockhorn, T., English, Z., Lacey, K., Goodwin, A., Marek, Y., Rombach, R.: Scaling Rectified Flow Transformers for High-Resolution Image Synthesis (2024). <https://arxiv.org/abs/2403.03206>
- [16] Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis (2023). <https://arxiv.org/abs/2307.01952>
- [17] OpenAI: DALL-E 3. OpenAI Blog (2023). <https://openai.com/dall-e>
- [18] Artificial Analysis: Text-to-Image Arena Leaderboard. Accessed: 2025-02-21 (2025). <https://artificialanalysis.ai/text-to-image/arena?tab=Leaderboard>
- [19] Innovating with generative ai: A hands-on comfyui workshop. In: Proceedings of the Special Interest Group on Computer Graphics and Interactive Techniques Conference Courses (2025). <https://doi.org/10.1145/3721241.3734004> . <https://dl.acm.org/doi/10.1145/3721241.3734004>
- [20] Royal College of Surgeons of England: Women in Surgery Statistics. <https://www.rcseng.ac.uk/careers-in-surgery/women-in-surgery/statistics/> (2022)