

Plug-and-Play Interpretable Responsible Text-to-Image Generation via Dual-Space Multi-facet Concept Control

Supplementary Material

This document provides detailed supplementary information to further explain the proposed framework and achieved results. It includes explanation of the training setup, algorithms, and technical specifics of the text embedding and diffusion latent manipulations, which form the core of our dual-space technique. It extends the analysis with further experiments, ablation studies, and implementation guidelines to ensure reproducibility. Our code can be accessed at <https://basim-azam.github.io/responsibleddiffusion/>

A. Further Algorithmic and Training Details

We explore the proposed technique at two intervention points in text-to-image generative pipeline; namely, at the text embedding space and the diffusion latent space. Below, we provide further training details of our modules used for each of these points.

Training RICE Module: Algorithm 1 presents an overview of the training procedure of RICE module, which distills the teacher encoder module $\mathcal{E}(\cdot)$ into the student encoder module $\mathcal{E}_{resp}(\cdot)$, to produce $z_{clip}^{distill}$. The underlying model for the RICE module ($\psi_{\mathcal{E}}$) is built upon the CLIP text encoder model. This encoder processes tokenized text input into embeddings of shape 77×768 , where the former represents the maximum sequence length and the latter is the dimensionality of each token’s embedding. The architecture of our model comprises a 12-layer Transformer with 8 attention heads per layer, facilitating the generation of contextualized token embeddings. To align these embeddings with the responsible concept space \mathcal{A}_{resp} , the module incorporates distillation process, where the pre-trained CLIP model serves as the teacher. A knowledge distillation loss $\mathcal{L}_{KD-clip}$ minimizes the discrepancy between the teacher’s embeddings and those of the student model. Additionally, a whitening transformation generates z^{zca} , combined with $z_{clip}^{distill}$ using a weighting factor α , yielding z_{clip}^{resp} . Backpropagation optimizes $\psi_{\mathcal{E}}$ iteratively across epochs and batches. This design ensure that the $\psi_{\mathcal{E}}$ integrates seamlessly in the T2I pipeline as a plug-in module for the generation of responsible images.

Training RIIDL Module: Algorithm 2 describes the training process for the RIIDL module with the underlying model $\psi_{\mathcal{D}}$ whose architecture is inspired by the U-Net structure of the Stable Diffusion v1.4. Our U-Net model consists of a downsampling path with 4 stages, each including 3×3 convolutional layers, group normalization, Swish activation, and residual connections. Downsampling reduces spatial resolution while progressively increasing fea-

ture channels, starting from 320 and doubling to 640, 1280, and 1280 across the stages. The bottleneck incorporates multi-head attention and residual blocks to capture global context and integrate cross-attention with text-conditioning inputs from the text encoder $\mathcal{E}(\cdot)$. The upsampling path mirrors the downsampling structure, with transpose convolutions or bilinear upsampling to restore spatial resolution and skip connections to incorporate features from corresponding downsampling stages.

During training, the RIIDL module aligns the student’s latents $z_{unet}^{distill,\tau}$ with the teacher’s latents $z_{unet}^{T,\tau}$ using a weighted distillation loss $\mathcal{L}_{KD-unet}$. Intermediate latents are decorrelated through a whitening transformation, producing $z_{unet}^{zca,\tau}$. These transformed latents are then combined with $z_{unet}^{distill,\tau}$ using a weighting factor β to generate the composite latent $z_{unet}^{resp,\tau}$. The final refined latent aligns with \mathcal{A}_{resp} while preserving high-quality image generation. The RIIDL module’s design ensures responsibility-driven latent adjustments without altering the original U-Net’s architecture, i.e adaptation of $\Psi_{\mathcal{D}}$ as a plug-in module to produce responsible latents.

Dual Space Inference: As we utilize both intervention points, i.e., text embedding space and the diffusion latent space, Algorithm 3 summarizes our dual-space inference process, harmonizing image generation while enforcing responsible aspects. The process combines RICE ($\psi_{\mathcal{E}}$) refined

Algorithm 1 Responsible and Interpretable CLIP Embeddings

Input: Text Encoder $\mathcal{E}(\cdot)$, Concept Set A_X , Text Prompts \mathcal{T} , Batch Size B , Number of Epochs N

Output: Trained RICE Module $\psi_{\mathcal{E}}$

- 1: Initialize $\psi_{\mathcal{E}}$ and optimizer
- 2: **for** epoch = 1 to N **do**
- 3: **for** each batch $B \subset \mathcal{T}$ **do**
- 4: $z_{clip}^T = \mathcal{E}(B)$ % Teacher Embeddings
- 5: $z_{clip}^{distill} = \psi_{\mathcal{E}}(B)$ % Student Embeddings
- 6:

$$\mathcal{L}_{KD-clip} = \frac{1}{|B|} \sum_{k=1}^{|B|} \|z_{clip,k}^T - z_{clip,k}^{distill}\|^2$$

- 7:
- 8: $z^{zca} = \mathcal{W}(z_{clip}^{distill} - \mu)$, $\mu = \frac{1}{|B|} \sum_{k \in B} z_{clip,k}^{distill}$

- 8:
- 9: $z_{clip}^{resp} = \alpha z_{clip}^{distill} + (1 - \alpha) z_{clip}^{zca}$

- 9: Update $\psi_{\mathcal{E}}$

- 10: **Return:** Trained $\psi_{\mathcal{E}}$
-

text embeddings and RIIDL ($\psi_{\mathcal{D}}$) enhanced latents. Responsible embeddings z_{clip}^{final} and latents z_{unet}^{final} are fused into a unified representation using control weights $\lambda_{\mathcal{E}}$ and $\lambda_{\mathcal{D}}$ such that $\lambda_{\mathcal{E}} + \lambda_{\mathcal{D}} = 1$. The final image \hat{I} adheres to \mathcal{A}_{resp} while preserving the base T2I model’s visual quality. The modular design keeps the pre-trained diffusion model unchanged, with RICE and RIIDL modules acting as external plug-ins in the T2I pipeline.

Hyperparameter Settings: The hyperparameters α in RICE and the corresponding parameter β in RIIDL balance semantic retention with responsibility aspects. They govern control between embeddings or latents and their decorrelated counterparts, subtly reinforcing responsible alignments while maintaining semantic and visual precision. The values of α and β are selected to give minimal but meaningful weight to the decorrelated components (i.e., 0.1). This ensures the influence without dominating the primary embeddings or latents. This value is empirically determined by inspecting visual quality of the outputs. During inference, a dominant emphasis on text embeddings is ensured through $\lambda_{\mathcal{E}} = 95\%$ for retaining the original semantics, while latent-space refinements via $\lambda_{\mathcal{D}} = 5\%$ is chosen to provide subtle enhancements that result in high image quality. All hyperparameters $\alpha, \beta, \lambda_{\mathcal{E}}$, and $\lambda_{\mathcal{D}}$ are determined through iterative experimentation and empirical analysis, balancing their adaptation across diverse scenarios. This balance enables effective module integration for RICE and RIIDL as plug-ins in the T2I pipeline for responsible image generation.

Training Setup: In our experiments, we use Stable Diffusion v1.4 checkpoint as the base model with a guidance scale of 7.5 for high-quality image generation. Concept spaces \mathcal{A}_{race} , \mathcal{A}_{age} , and \mathcal{A}_{gender} are trained using 1,000 synthesized images per concept, optimized over 10,000 steps. During inference, these concept spaces are applied within the dual-space control framework to enforce adherence to \mathcal{A}_{resp} . We empirically find that that higher values of $\lambda_{\mathcal{D}}$ yield better results, emphasizing the importance of latent space interventions in maintaining structural fidelity while ensuring responsibility constraints. The experiments are carried out on a single node of the an HPC system equipped with 4 NVIDIA A100 GPUs (80GB each), 495GB of RAM, and 32 CPU cores.

B. Further Results

In the main paper, we analyzed the dual-space control framework’s effectiveness in mitigating biases across professions using WinoBias dataset [7], as detailed in sections 5 and 6 of the paper. To extend the analysis we provide additional results to demonstrate further capabilities of our framework here.

Figures B.1 and B.2 provide further examples of successful control over responsible concepts using our tech-

Algorithm 2 Responsible Interpretable Intermediate Diffusion Latents

Input: Diffusion Model $\mathcal{D}(\cdot)$, Concept Set \mathcal{A}_X , Batch Size B , Number of Epochs N

Output: Trained RIIDL Module $\psi_{\mathcal{D}}$

- 1: Initialize $\psi_{\mathcal{D}}$ and optimizer
 - 2: **for** epoch = 1 to N **do**
 - 3: **for** each batch B **do**
 - 4: Uniformly sample time step $\tau \sim U[0, 1]$
 - 5: $z_{unet}^{T,\tau} = \mathcal{D}(\mathcal{E}(\cdot))$ % Teacher Latents
 - 6: $z_{unet}^{distill,\tau} = \psi_{\mathcal{D}}(B)$ % Student Latents
 - 7: $L_{KD-unet} = \omega(\lambda_{\tau}) \|z_{unet}^{T,\tau} - z_{unet}^{distill,\tau}\|^2$
 - 8: $z_{unet}^{zca,\tau} = W(z_{unet}^{distill,\tau} - \mu)$
 - 9: $z_{unet}^{resp,\tau} = \beta z_{unet}^{distill,\tau} + (1 - \beta) z_{unet}^{zca,\tau}$
 - 10: Update $\psi_{\mathcal{D}}$
 - 11: **Return:** Trained $\psi_{\mathcal{D}}$
-

Algorithm 3 Dual-Space Integration and Inference Process

Input: Text Prompt t , RICE Module $\psi_{\mathcal{E}}$, RIIDL Module $\psi_{\mathcal{D}}$, T2I Model Ψ , Weights $\lambda_{\mathcal{E}}, \lambda_{\mathcal{D}}$

Output: Responsible Image \hat{I}

- 1: Compute original text embeddings:

$$z_{clip} = \mathcal{E}(t)$$

- 2: Compute responsible embeddings using RICE:

$$z_{clip}^{final} = z_{clip} + \sum_k \gamma_k z_{k,clip}^{resp}$$

- 3: Define the responsible embedding module:

$$\Psi_{\mathcal{E}} = \mathcal{E}_{resp}(t), \quad \text{where } \mathcal{E}_{resp} \text{ uses } z_{clip}^{final}$$

- 4: Compute original latent representation using the diffusion model:

$$z_{unet} = \mathcal{D}(\mathcal{E}(t))$$

- 5: Compute responsible latents using RIIDL:

$$z_{unet}^{final} = z_{unet} + \sum_k \gamma_k z_{k,unet}^{resp}$$

- 6: Define the responsible latent module:

$$\Psi_{\mathcal{D}} = \mathcal{D}_{resp}(\mathcal{E}(t)), \quad \text{where } \mathcal{D}_{resp} \text{ uses } z_{unet}^{final}$$

- 7: Combine embedding and latent representations via dual-space control:

$$\Psi_{resp}(t) = \lambda_{\mathcal{E}} \Psi_{\mathcal{E}} + \lambda_{\mathcal{D}} \Psi_{\mathcal{D}}, \quad \lambda_{\mathcal{E}} + \lambda_{\mathcal{D}} = 1$$

- 8: Generate the final responsible image:

$$\hat{I} = \Psi_{resp}(t)$$

- 9: **Return:** Responsible Image \hat{I}
-

nique. Figure B.1 provides a pairwise comparison of baseline Stable Diffusion (SD) and our responsible generation outputs, where baseline outputs often reflect stereotypical associations. By incorporating our technique, the genera-

tion aligns with the responsible concept space \mathcal{A}_X . This alignment ensures balanced and non-stereotypical representations across the professions. Figure B.2 presents the ability of the proposed approach to incorporate different responsible attributes of age, gender and race in the generated content for different professions. Figure B.3 extends this analysis by incorporating multiple aspects from different concepts simultaneously to generate responsible outputs. It is observable that the generated outputs remain high-quality and they preserve the original semantics while incorporating the multi-facet responsible concepts that we like to be considered in the outputs. This is a key strength of our method.

In section 6 of the main paper, we presented curtailed results about the transitions caused by RICE and RIIDL modules. Figures B.4 and B.5 provide extended results for the same purpose. The results emphasize the roles of RICE and RIIDL modules in enabling fine-grained control. Figure B.4 demonstrates smooth interpolation within the embedding space, transitioning across attributes such as gender and race, with corner anchor images representing distinct combinations. Figure B.5 highlights RIIDL’s ability to refine latent space representations during diffusion, aligning outputs with responsible attributes at various denoising steps. Additional examples in Figures B.6a and B.6b exhibit the further scalability of our framework. Figure B.6a integrates additional attributes (e.g., smile, glasses) into professional contexts, while Figure B.6b extends the approach to abstract prompts such as animals and activities (e.g., “a dog skateboarding”). These results validate the versatility and adaptability of the proposed method across diverse cases.

In the main paper, we presented quantitative results on Winobias dataset to quantify bias and debiasing for randomly selected classes (due to space restrictions). Here, Tables B.1 and B.2 present extensive results along the same lines to demonstrate the efficacy of the proposed framework against prior works. Table B.1 reports deviation ratio values (Δ) across gender and racial attributes, showing that our method achieves the lowest average deviation across all attributes, outperforming standard Stable Diffusion and competing methods such as Unified Concept Editing (U) [2] and Vector Interpret Diffusion (V) [3]. Table B.2 further corroborates these findings by presenting debiasing performance results, where our approach achieves the lowest average deviation, indicating superior fairness.

Supplement References

- [1] Ching-Yao Chuang et al. “Debiasing vision-language models via biased prompts”. In: *arXiv preprint arXiv:2302.00070* (2023).
- [2] Rohit Gandikota et al. “Unified concept editing in diffusion models”. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*. 2024, pp. 5111–5120.
- [3] Hang Li et al. “Self-discovering interpretable diffusion latent directions for responsible text-to-image generation”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. June 2024.
- [4] Hadas Orgad, Bahjat Kawar, and Yonatan Belinkov. “Editing implicit assumptions in text-to-image diffusion models”. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2023, pp. 7053–7061.
- [5] Robin Rombach et al. “High-resolution image synthesis with latent diffusion models”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2022, pp. 10684–10695.
- [6] Zihao Wang et al. “Concept algebra for text-controlled vision models”. In: *arXiv preprint arXiv:2302.03693* 2 (2023).
- [7] Jieyu Zhao et al. “Gender bias in coreference resolution: Evaluation and debiasing methods”. In: *arXiv preprint arXiv:1804.06876* (2018).



Figure B.1. Pairwise Comparison of Baseline and Responsible Generation Outputs. Each pair compares Stable Diffusion (SD) outputs (top) with our responsible generation plugin (bottom) across age, gender, and race attributes for each profession in Winobias [7] dataset. SD outputs often show stereotypical biases, such as gender or racial associations among the professions. Using our technique, incorporating ψ_E and ψ_D , these biases are mitigated by aligning outputs with responsible aspects. Our method ensures fair, diverse, and high-quality generation while removing harmful stereotypes.



Figure B.2. Controlled responsible generation by incorporating different individual concepts in the outputs. The figure showcases images generated by incorporating our technique into Stable Diffusion with individual control over attributes such as age, gender, and race. The prompt "A <profession-name>" is used to generate the image, and individual target concept space from (\mathcal{A}_X) is plugged in using the proposed approach. Each row highlights selected profession from the Winobias dataset [7], while columns represent the incorporation of concept spaces.



Figure B.3. Responsible generation by incorporating multiple aspects simultaneously. The figure highlights the ability of the proposed approach to handle composite attributes across professions by leveraging dual-space control. Each row corresponds to a specific profession, with columns showing combination of attributes for age, gender, and race, incorporated for responsible generation.



Figure B.4. Text embedding space transition caused via RICE. The figure showcases progressive transition by manipulating $\mathcal{A}_{\text{gender}}$ and $\mathcal{A}_{\text{race}}$. The corner anchor images represent distinct combinations; images resulting from the interpolated intermediate embeddings present the continuous transition spectrum and showcase the changes as the generated outputs align with the concepts in \mathcal{A}_X .



Figure B.5. Latent-space refinement caused by RIIDL. The figure showcases RIIDL module modulating latent diffusion space, progressively aligning outputs with the responsible attributes \mathcal{A}_x . Starting from the latent state for <a doctor>, RIIDL enforces control over $\mathcal{A}_{\text{gender}}$ and $\mathcal{A}_{\text{race}}$ by injecting adjustments during denoising.



(a)



(b)

Figure B.6. Extended attribute control with the proposed method. Additional examples of extending \mathcal{A}_X with attributes beyond core responsible concepts are presented. (a) Shows randomly selected four professions (scientists, lawyers, teachers, and nurses) for which images are enhanced seamlessly by integrating attributes such as A_{smile} and $A_{glasses}$ alongside age, gender, and race controls. Notices that the underlying concepts of race, gender and age remain effective besides the additional attributes of smile and glasses. (b) Highlights extensions in a similar manner for animal and object prompts ("a dog," "a cat," "a squirrel", and "a boy") with attributes such as A_{eating} , $A_{skateboarding}$, $A_{jumping}$, showcasing adaptability and scalability. These results validate the ability of the method to incorporate additional concept spaces while maintaining semantic coherence and visual quality.

Table B.1. Deviation ratio values across gender and racial attributes in the WinoBias dataset. Lower values indicate better performance. Results for methods: SD, U, and V are also included, where SD: Original Standard Diffusion Model, U: Unified Concept Editing [2], V: Vector Interpret Diffusion [3].

Attribute	Gender(↓)				Gender-Pro(↓)				Race(↓)				Race-Pro(↓)			
	SD [5]	U [2]	V [3]	Ours	SD [5]	U [2]	V [3]	Ours	SD	U [2]	V [3]	Ours	SD [5]	U [2]	V [3]	Ours
Analyst	0.70	0.20	0.04	0.05	0.54	0.04	0.08	0.08	0.82	0.29	0.24	0.07	0.77	0.20	0.41	0.10
Assistant	0.02	0.14	0.37	0.04	0.48	0.80	0.45	0.09	0.38	0.17	0.24	0.06	0.24	0.26	0.12	0.08
Attendant	0.16	0.09	0.24	0.02	0.78	0.08	0.26	0.12	0.37	0.16	0.22	0.07	0.67	0.37	0.13	0.10
Baker	0.82	0.29	0.07	0.05	0.64	1.00	0.09	0.11	0.83	0.14	0.12	0.06	0.72	0.32	0.16	0.09
CEO	0.92	0.28	0.15	0.04	0.90	0.58	0.21	0.13	0.38	0.13	0.22	0.08	0.31	0.08	0.22	0.12
Carpenter	0.92	0.06	0.02	0.06	1.00	1.00	0.03	0.14	0.91	0.12	0.28	0.08	0.83	0.65	0.26	0.11
Cashier	0.74	0.16	0.63	0.03	0.92	0.92	0.66	0.15	0.45	0.43	0.34	0.07	0.46	0.41	0.30	0.09
Cleaner	0.54	0.33	0.06	0.04	0.30	0.80	0.08	0.09	0.10	0.28	0.14	0.06	0.45	0.55	0.26	0.08
Clerk	0.14	0.23	0.42	0.04	0.58	0.96	0.46	0.12	0.46	0.25	0.16	0.07	0.59	0.38	0.16	0.10
Const. Worker	1.00	0.06	0.80	0.05	1.00	0.24	0.84	0.11	0.41	0.16	0.26	0.08	0.44	0.29	0.25	0.12
Cook	0.72	0.03	0.04	0.05	0.02	0.36	0.07	0.08	0.56	0.15	0.30	0.07	0.18	0.49	0.14	0.10
Counselor	0.00	0.40	0.45	0.04	0.56	1.00	0.47	0.16	0.72	0.19	0.16	0.07	0.36	0.79	0.12	0.09
Designer	0.12	0.07	0.28	0.03	0.72	0.84	0.30	0.10	0.14	0.23	0.10	0.06	0.18	0.34	0.15	0.11
Developer	0.90	0.51	0.54	0.04	0.92	0.96	0.58	0.14	0.41	0.23	0.30	0.08	0.32	0.20	0.39	0.13
Doctor	0.92	0.20	0.08	0.05	0.52	0.32	0.10	0.10	0.92	0.07	0.26	0.07	0.59	0.52	0.15	0.12
Driver	0.90	0.21	0.12	0.04	0.48	0.60	0.14	0.09	0.34	0.23	0.16	0.08	0.25	0.26	0.07	0.10
Farmer	1.00	0.41	0.21	0.03	0.98	0.12	0.30	0.08	0.95	0.27	0.50	0.07	0.39	0.82	0.28	0.09
Guard	0.78	0.12	0.14	0.04	0.76	0.08	0.17	0.15	0.20	0.16	0.12	0.09	0.35	0.23	0.14	0.10
Hairdresser	0.92	0.16	0.84	0.05	0.88	0.46	0.86	0.17	0.45	0.31	0.42	0.07	0.38	0.05	0.23	0.12
Housekeeper	0.96	0.41	0.90	0.04	1.00	1.00	0.91	0.13	0.45	0.07	0.28	0.07	0.45	0.41	0.34	0.09
Janitor	0.96	0.16	0.44	0.05	0.94	0.08	0.46	0.14	0.35	0.14	0.24	0.06	0.40	0.24	0.07	0.10
Laborer	1.00	0.09	0.41	0.05	0.98	0.08	0.43	0.08	0.33	0.40	0.24	0.10	0.53	0.38	0.20	0.11
Lawyer	0.68	0.30	0.10	0.04	0.36	0.18	0.17	0.12	0.64	0.20	0.18	0.08	0.52	0.14	0.13	0.09
Librarian	0.66	0.07	0.81	0.03	0.54	0.40	0.86	0.09	0.85	0.28	0.42	0.06	0.74	0.16	0.27	0.10
Manager	0.46	0.19	0.10	0.05	0.62	0.40	0.91	0.10	0.69	0.17	0.24	0.08	0.41	0.17	0.19	0.12
Mechanic	1.00	0.23	0.34	0.04	0.98	0.48	0.46	0.11	0.64	0.22	0.14	0.07	0.47	0.44	0.05	0.08
Nurse	1.00	0.39	0.96	0.04	0.98	0.84	0.43	0.18	0.76	0.25	0.30	0.07	0.39	0.79	0.08	0.14
Physician	0.78	0.42	0.06	0.05	0.30	0.16	0.17	0.08	0.67	0.08	0.18	0.08	0.46	0.58	0.02	0.10
Receptionist	0.84	0.38	0.88	0.04	0.98	0.96	0.86	0.17	0.88	0.10	0.36	0.07	0.74	0.14	0.25	0.11
Salesperson	0.68	0.38	0.04	0.04	0.54	0.12	0.18	0.11	0.69	0.32	0.26	0.07	0.66	0.19	0.36	0.13
Secretary	0.64	0.10	0.76	0.04	0.92	0.96	0.44	0.14	0.37	0.35	0.24	0.06	0.55	0.43	0.32	0.09
Sheriff	1.00	0.10	0.24	0.05	0.98	0.24	0.97	0.12	0.82	0.17	0.18	0.08	0.74	0.35	0.27	0.10
Supervisor	0.64	0.26	0.16	0.05	0.52	0.46	0.12	0.10	0.49	0.14	0.14	0.07	0.45	0.31	0.14	0.12
Tailor	0.56	0.27	0.14	0.04	0.78	0.48	0.89	0.08	0.16	0.20	0.10	0.07	0.14	0.19	0.13	0.09
Teacher	0.30	0.06	0.51	0.04	0.48	0.16	0.07	0.13	0.51	0.10	0.04	0.08	0.26	0.23	0.21	0.10
Writer	0.04	0.31	0.49	0.04	0.26	0.52	0.91	0.10	0.86	0.23	0.26	0.06	0.69	0.38	0.07	0.09
Average	0.68	0.22	0.36	0.04	0.70	0.52	0.61	0.14	0.56	0.21	0.23	0.07	0.48	0.35	0.20	0.10

Table B.2. Debiasing performance across professions in Winobias dataset [7]. The metric $\Delta = 0$ indicate ideal debiasing. Our method shows very strong performance compared to previous approaches.

Profession	SD	CA [6]	DVL [1]	TIME [4]	TIMEP [2]	Ours
Attendant	0.13 ± 0.06	0.23 ± 0.08	0.30 ± 0.04	0.50 ± 0.01	0.38 ± 0.11	0.09 ± 0.04
Cashier	0.67 ± 0.04	0.71 ± 0.10	0.23 ± 0.07	0.46 ± 0.01	0.23 ± 0.15	0.16 ± 0.06
Teacher	0.42 ± 0.01	0.46 ± 0.00	0.11 ± 0.05	0.34 ± 0.06	0.07 ± 0.06	0.05 ± 0.01
Nurse	0.99 ± 0.01	0.91 ± 0.05	0.87 ± 0.01	0.34 ± 0.03	0.20 ± 0.09	0.20 ± 0.07
Assistant	0.19 ± 0.05	0.20 ± 0.07	0.35 ± 0.15	0.32 ± 0.06	0.57 ± 0.08	0.14 ± 0.06
Secretary	0.88 ± 0.01	0.65 ± 0.07	0.65 ± 0.01	0.58 ± 0.09	0.71 ± 0.02	0.10 ± 0.10
Cleaner	0.38 ± 0.04	0.11 ± 0.06	0.18 ± 0.04	0.58 ± 0.07	0.79 ± 0.04	0.13 ± 0.07
Receptionist	0.99 ± 0.01	0.90 ± 0.08	0.74 ± 0.04	0.36 ± 0.10	0.24 ± 0.12	0.10 ± 0.01
Clerk	0.10 ± 0.07	0.11 ± 0.08	0.10 ± 0.04	0.40 ± 0.03	0.76 ± 0.05	0.23 ± 0.06
Counselor	0.06 ± 0.05	0.30 ± 0.03	0.10 ± 0.07	0.74 ± 0.08	0.41 ± 0.06	0.11 ± 0.02
Designer	0.23 ± 0.05	0.25 ± 0.12	0.48 ± 0.06	0.44 ± 0.06	0.23 ± 0.16	0.07 ± 0.05
Hairdresser	0.74 ± 0.11	0.37 ± 0.16	0.61 ± 0.04	0.32 ± 0.01	0.41 ± 0.09	0.10 ± 0.04
Writer	0.15 ± 0.03	0.07 ± 0.03	0.45 ± 0.04	0.54 ± 0.08	0.52 ± 0.08	0.09 ± 0.08
Housekeeper	0.93 ± 0.04	0.68 ± 0.18	0.80 ± 0.07	0.32 ± 0.03	0.68 ± 0.07	0.08 ± 0.05
Baker	0.81 ± 0.01	0.19 ± 0.04	0.72 ± 0.05	0.40 ± 0.04	0.19 ± 0.14	0.20 ± 0.08
Librarian	0.86 ± 0.06	0.66 ± 0.07	0.34 ± 0.06	0.26 ± 0.05	0.35 ± 0.01	0.04 ± 0.02
Tailor	0.30 ± 0.01	0.21 ± 0.05	0.33 ± 0.11	0.50 ± 0.03	0.03 ± 0.00	0.20 ± 0.01
Driver	0.97 ± 0.02	0.20 ± 0.07	0.65 ± 0.04	0.48 ± 0.09	0.11 ± 0.08	0.21 ± 0.07
Supervisor	0.50 ± 0.01	0.07 ± 0.03	0.43 ± 0.04	0.50 ± 0.07	0.42 ± 0.03	0.26 ± 0.04
Janitor	0.91 ± 0.05	0.71 ± 0.06	0.75 ± 0.05	0.36 ± 0.08	0.47 ± 0.12	0.16 ± 0.04
Cook	0.82 ± 0.04	0.48 ± 0.16	0.52 ± 0.07	0.38 ± 0.03	0.15 ± 0.10	0.03 ± 0.02
Laborer	0.99 ± 0.01	0.81 ± 0.06	0.98 ± 0.03	0.48 ± 0.08	0.24 ± 0.09	0.09 ± 0.02
Constr. worker	1.00 ± 0.00	0.95 ± 0.01	1.00 ± 0.00	0.40 ± 0.01	0.15 ± 0.05	0.06 ± 0.04
Developer	0.90 ± 0.03	0.74 ± 0.02	0.90 ± 0.04	0.50 ± 0.01	0.47 ± 0.07	0.08 ± 0.02
Carpenter	0.92 ± 0.05	0.84 ± 0.01	0.98 ± 0.01	0.52 ± 0.06	0.52 ± 0.05	0.06 ± 0.02
Manager	0.54 ± 0.06	0.15 ± 0.01	0.30 ± 0.05	0.38 ± 0.05	0.10 ± 0.01	0.19 ± 0.07
Lawyer	0.46 ± 0.08	0.13 ± 0.06	0.52 ± 0.05	0.64 ± 0.03	0.15 ± 0.03	0.30 ± 0.07
Farmer	0.97 ± 0.02	0.58 ± 0.09	0.97 ± 0.02	0.46 ± 0.02	0.27 ± 0.08	0.09 ± 0.01
Salesperson	0.60 ± 0.08	0.18 ± 0.05	0.07 ± 0.05	0.52 ± 0.05	0.05 ± 0.01	0.11 ± 0.05
Physician	0.62 ± 0.14	0.36 ± 0.10	0.70 ± 0.07	0.56 ± 0.06	0.49 ± 0.04	0.11 ± 0.04
Guard	0.86 ± 0.02	0.43 ± 0.12	0.48 ± 0.06	0.30 ± 0.10	0.10 ± 0.12	0.12 ± 0.07
Analyst	0.58 ± 0.12	0.24 ± 0.18	0.71 ± 0.02	0.52 ± 0.03	0.13 ± 0.05	0.07 ± 0.02
Mechanic	0.99 ± 0.01	0.65 ± 0.04	0.92 ± 0.01	0.38 ± 0.09	0.21 ± 0.04	0.13 ± 0.08
Sheriff	0.99 ± 0.01	0.38 ± 0.22	0.82 ± 0.08	0.22 ± 0.05	0.10 ± 0.05	0.06 ± 0.03
CEO	0.87 ± 0.03	0.25 ± 0.11	0.37 ± 0.11	0.28 ± 0.04	0.18 ± 0.05	0.07 ± 0.03
Doctor	0.78 ± 0.04	0.40 ± 0.02	0.50 ± 0.04	0.58 ± 0.03	0.41 ± 0.08	0.06 ± 0.01