# Zero-shot Conditional Face Generation using Medical Knowledge Graphs

## 1 General Info

**Project Title**: Zero-shot Conditional Face Generation using Medical Knowledge Graphs
**Supervisors**: Yousef Yeganeh, Azade Farshad
**Contact Email**: y.yeganeh@tum.de, azade.farshad@tum.de

## 2 Background and Motivation

Facial image generation has numerous applications in computer vision, graphics, as well as biomedicine. Moreover, generating facial images with specific disease categories is even more challenging due to the scarcity and privacy of training data [5]. However, such images could have a huge impact on improving healthcare in everyday life by enabling early diagnosis, disease monitoring, and patient education [8]. Therefore, there is a need for novel approaches that can leverage the existing knowledge to generate high-quality facial images with distinct disease categories in a zero-shot manner. One possible source of such knowledge is medical knowledge graphs (KGs) [1], which are structured representations of biomedical concepts and their relations. Medical KGs can provide rich and accurate information about the symptoms, causes, and treatments of different diseases, as well as the facial attributes that are associated with them. For example, a medical KG can tell us that Down syndrome is a genetic disorder that is likely to cause a flattened appearance to the face, almond-shaped eyes, and a short neck.

## 3 Project Abstract

The goal of this project is to generate realistic faces with various disease categories using generative models. In this project, we propose to first generate disease profiles based on medical knowledge graphs [1] and large language models [9]. These profiles would describe the facial attributes of people with different diseases at different ages. Next, we would use existing generative models such as Stable Diffusion [7] and condition them on the disease profiles. Finally, we would produce longitudinal facial images of the diseases at different ages [6]. The generated images would be evaluated against real-world face disease datasets [4, 2, 3] and using disease classification models. The project consists of the following steps: 1) Familiarization with the literature, 2) Training a supervised generative model on the DSF dataset as baseline, 3) Extracting disease profiles from the medical knowledge graphs, 4) Zero-shot generation of disease-specific faces from the disease profiles, 5) Evaluation and documentation.

## 4 Technical Prerequisites

- Good background in machine learning and deep learning
- Experience in PyTorch
- Experience in Python
- Experience with Generative Models

## 5 Benefits

- Weekly supervision and discussions
- Possible novelty of the research
- The results of this work are intended to be published in a conference or journal

## 6 Work packages and Time-plan

* The dates are adopted from the previous year and are not finalized yet.

| | Description | # Students | From | To |
|---|---|---|---|---|
| WP1 | Familiarizing with the literature. | 4 | 24.10 | 31.10 |
| WP2 | Implementing the baselines | 4 | 31.10 | 14.11 |
| WP3 | Improving the baselines and validation on relevant datasets | 4 | 14.11 | 27.11 |
| | Midterm Presentation (Date is not finalized) | 4 | 27.11 | 05.12 |
| WP4 | Implementing the model | 4 | 05.12 | 19.12 |
| WP5 | Finalizing the results and evaluation | 4 | 19.12 | 07.02 |
| | Final Presentation (Date is not finalized) | 4 | 07.02 | 14.02 |

Table 1: Project Timeline

## References

[1] Payal Chandak, Kexin Huang, and Marinka Zitnik. Building a knowledge graph to enable precision medicine. *Nature Scientific Data*, 2023.

[2] Bo Jin. Disease-specific faces, 2020.

[3] Bo Jin. Disease-specific faces 2, 2023.

[4] Bo Jin, Leandro Cruz, and Nuno Gonçalves. Deep facial diagnosis: Deep transfer learning from face recognition to facial diagnosis. *IEEE Access*, 8:123649–123661, 2020.

[5] Anwesha Mohanty, Alistair Sutherland, Marija Bezbradica, and Hossein Javidnia. High fidelity synthetic face generation for rosacea skin condition from limited data. *arXiv preprint arXiv:2303.04839*, 2023.

[6] Chi Nhan Duong, Khoa Luu, Kha Gia Quach, and Tien D Bui. Longitudinal face modeling via temporal deep restricted boltzmann machines. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5772–5780, 2016.

[7] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.

[8] Kuan Wang and Jiebo Luo. Detecting visually observable disease symptoms from faces. *EURASIP Journal on Bioinformatics and Systems Biology*, 2016:1–8, 2016.

[9] Michihiro Yasunaga, Antoine Bosselut, Hongyu Ren, Xikun Zhang, Christopher D Manning, Percy S Liang, and Jure Leskovec. Deep bidirectional language-knowledge graph pretraining. *Advances in Neural Information Processing Systems*, 35:37309–37323, 2022.