Dr. George Close — ML Researcher in Speech & Audio

I am a machine learning researcher specialising in speech and audio. I am currently a Member of Technical Staff on the Audio team at Zyphra, where I help build large-scale text-to-speech and audio foundation models — including the ZONOS family of TTS systems — working across model architecture, data pipelines and evaluation.

I take ideas from research all the way to production: training and shipping models at scale on large GPU clusters, and grounding them in how humans actually perceive sound. This builds on my PhD in Computer Science from the University of Sheffield (perceptually-motivated speech enhancement and quality assessment) and a first-class BSc from Cardiff University, with published work spanning speech enhancement, speech quality prediction and ASR.

I am especially interested in the place where large speech models meet human perception — building systems that don't just score well on metrics, but genuinely sound right to people. Always happy to talk speech, audio and TTS: george@zyphra.com.

Interests

- Text To Speech (TTS) & speech synthesis
- Audio foundation models
- Self-Supervised Speech Representations
- Speech Enhancement / Noise Reduction
- Speech Quality / Intelligibility machine perception and prediction
- Automatic Speech Recognition (ASR)
- Human perception of digital audio
- Neural systems for hearing aids & edge devices
- Deepfake detection / adversarial attacks

Technical Skills

- Large-scale model training on multi-GPU clusters
- Python — PyTorch, SpeechBrain, NVIDIA NeMo, SciPy
- HuggingFace ecosystem / OpenAI API
- Distributed training & data pipelines for audio
- Linux / Bash scripting
- Git / GitHub / GitLab
- C++, Java, SQL, MATLAB

Experience & Qualifications

1. Sep 2025 — Present **Member of Technical Staff @ Zyphra**
   San Francisco, CA, USA Research and development on ZONOS, a large-scale text-to-speech model - spanning model architecture, data pipelines and evaluation.

2. Nov 2024 — Sep 2025 **Speech Data Scientist @ ConnexAI**
   Manchester, UK Built and deployed production speech-processing systems, including data filtering for in-the-wild speech corpora and non-intrusive speech quality prediction.
3. May 2024 — Aug 2024 **Yamaha Research and Development (Internship)**
   Hamamatsu, Japan Research internship applying machine learning to audio and music signal processing within an industrial R&D team. Project focused on modeling human perception of music spatiality.
4. Oct 2020 — Jan 2025 **PhD Computer Science + Graduate Teaching Assistant**
   University of Sheffield, UK
   Thesis: *Perceptually Motivated Speech Enhancement* Researched neural speech enhancement guided by human perception — using speech quality metrics and self-supervised representations as loss functions. Authored 10+ first-author papers and taught undergraduate courses as a GTA.
5. Aug 2017 — Aug 2020 **BSc Computer Science (First Class Honours)**
   Cardiff University, UK
   Thesis: *Majel — Voice control for Command Line Interfaces* Graduated with First Class Honours, with a final-year project building a speech-driven interface for the command line.

Papers & Publications

I am an author on 18 papers, of which 10 I am first author. These have amassed 161 citations with an h-index of 8. Full list on Google Scholar.

- ### ZONOS2 Technical Report

  Gabriel Clark, Sofian Mejjoute, Mohamed Osman, George Close, Beren Millidge

  2026 Zyphra Technical Report

  arXiv

- ### WhiSQA: Non-Intrusive Speech Quality Prediction Using Whisper Encoder Features

  George Close, Kris Hong, Thomas Hain, Stefan Goetze

  2025 SPECOM 2025

  arXiv GitHub

- ### Whilter: A Whisper-based Data Filter for "In-the-Wild" Speech Corpora Using Utterance-level Multi-Task Classification

  William Ravenscroft, George Close, Kit Bower-Morris, Jamie Stacey, Dmitry Sityaev, Kris Y. Hong

  2025 Interspeech 2025

  arXiv Interspeech 2025 Dataset

- ### Hallucination in Perceptual Metric-Driven Speech Enhancement Networks

  George Close, Thomas Hain, Stefan Goetze

  2024 EUSIPCO 2024

  arXiv EUSIPCO 2024 Listening Test Examples

- ## Using Speech Foundational Models in Loss Functions for Hearing Aid Speech Enhancement

  Robert Sutherland, George Close, Thomas Hain, Stefan Goetze, Jon Barker

  2024 EUSIPCO 2024

  EUSIPCO 2024

- ## Transcription-Free Fine-Tuning of Speech Separation Models for Noisy and Reverberant Multi-Speaker Automatic Speech Recognition

  William Ravenscroft, George Close, Stefan Goetze, Thomas Hain, Mohammad Soleymanpour, Anurag Chowdhury, Mark C. Fuhs

  2024 Interspeech 2024

  arXiv Interspeech 2024

- ## Non-Intrusive Speech Intelligibility Prediction for Hearing-Impaired Users using Intermediate ASR Features and Human Memory Models

  Rhiannon Mogridge, George Close, Robert Sutherland, Thomas Hain, Jon Barker, Stefan Goetze, Anton Ragni

  2024 ICASSP 2024 🏆 2nd Place — Clarity Prediction Challenge 2

  Technical Report arXiv ICASSP 2024

- ## Multi-CMGAN+/+: Leveraging Multi-Objective Speech Quality Metric Prediction for Speech Enhancement

  George Close, Thomas Hain, Stefan Goetze

  2024 ICASSP 2024

  arXiv ICASSP 2024 GitHub

- ## CMGAN+/+: The University of Sheffield CHiME-7 UDASE Challenge Speech Enhancement System

  George Close, William Ravenscroft, Thomas Hain, Stefan Goetze

  2023 CHiME-7 UDASE Challenge 🏆 Challenge Entry

  Technical Report Results

- ## Non-Intrusive Intelligibility Predictor for Hearing-Impaired Individuals using Self-Supervised Speech Representations

  George Close, Thomas Hain, Stefan Goetze

  2023 SPARKS Workshop 2023

  arXiv SPARKS Workshop 2023

- ## The Effect of Spoken Language on Speech Enhancement using Self-Supervised Speech Representation Loss Functions

  George Close, Thomas Hain, Stefan Goetze

  2023 WASPAA 2023

  arXiv WASPAA 2023 GitHub

- ## Perceive and Predict: Self-Supervised Speech Representation Based Loss Functions for Speech Enhancement

  George Close, William Ravenscroft, Thomas Hain, Stefan Goetze

  2023 ICASSP 2023

  arXiv ICASSP 2023 Audio Examples

- ## PAMGAN+/-: Improving Phase-Aware Speech Enhancement Performance via Expanded Discriminator Training

  George Close, Thomas Hain, Stefan Goetze

  2023 154th AES Convention 🏆 Student Technical Paper Award

  AES e-Library

- ## Non-Intrusive Speech Intelligibility Metric Prediction for Hearing-Impaired Individuals — Clarity Prediction Challenge 1

  George Close, Samuel Hollands, Thomas Hain, Stefan Goetze

  2022 Interspeech 2022

  Interspeech 2022

- ## MetricGAN+/-: Increasing Robustness of Noise Reduction on Unseen Data

  George Close, Thomas Hain, Stefan Goetze

  2022 EUSIPCO 2022

  arXiv EUSIPCO 2022 Audio Examples

Talks & Presentations

AES [TC-MLAI] 2022 Talk — "Teaching AI to hear like we do: psychoacoustics in machine learning"