

**UTTER**

**Unified Transcription and Translation for
Extended Reality
(UTTER)**

**Horizon Europe Research and Innovation Action
Number: 101070631
D9.1 – UTTER First Ethics Review**

Nature	ETHICS	Work Package	WP9
Due Date	30/09/2023	Submission Date	30/09/2023
Main authors	Alexandra Birch (UEDIN)		
Co-authors	Wilker Aziz (UvA), André F. T. Martins (IT), José G. C. de Souza (UNB), Laurent Besacier (NAVER)		
Reviewers	Chrysoula Zerva (IT)		
Keywords	ethics, data		
Version Control			
v0.1	Status	Draft	21/09/2023
v1.0	Status	Final	23/09/2023



Contents

1	Introduction	3
2	UTTER ethics process	4
3	Protection of personal data	4
4	Experimental data	5
5	Broader ethical concerns	6
6	Conclusion	7

1 Introduction

The goal of UTTER is the provision of multilingual and multimodal (speech and text) intelligent assistant capabilities for online meetings and customer service support. We will process, analyse and distill conversations with the ultimate goal of improving communication between participants. These conversations can contain personal data, which means that the project needs to carefully address ethical issues relating to privacy. The objective of this workpackage is to ensure that we follow the principles of: Transparency, Accountability and Fairness.

We take the Data Ethics Framework ¹ prepared by the Horizon Europe as our guiding principles. The main guidance is to ensure that the public benefit of the project is balanced against the risks of pursuing the project. The six principles of the framework are:

- We define and understand public benefit and user need
- Involve diverse expertise, including ethical experts
- We are compliant with applicable UK, EU and national law
- That we implement our work responsibly, increasing user confidence in our tools and services
- Review the quality and limitations of the data
- Evaluate and consider wider policy implication

Our project will require four kinds of data: the use of existing multilingual and conversational data sets, the creation of annotated data for training and automatic testing our models, data that is processed in the use of our platforms, and data generated from field tests eg. surveys of users. Each type of data has different implications for data management and ethics.

This deliverable is aligned with the deliverable D1.1 (*Data Management Plan*) which was submitted on 31 March 2023. In this D1.1, we address the types of data used in the project, and data storage and reuse.

Here we outline the ethical review and recommendation process that we have established in the UTTER project, and we group the ethical issues which arise into these categories:

- Ethics process
- Protection of personal data;
- Experimental data;
- Broader ethical concerns including dual use (military and defence) implications of UTTER technologies, and the social impact of the tools and technologies developed by UTTER.

In this report we summarise the ethical issues raised.

¹ https://ec.europa.eu/info/funding-tenders/opportunities/docs/2021-2027/horizon/guidance/ethics-and-data-protection_he_en.pdf

2 UTTER ethics process

The project has an Ethics Committee which was formed at the kickoff meeting. It is comprised of the same members of the consortium who form part of the project board, one representative per project partner. The Ethics Committee meets once a year, when possible in person, to address any issues regarding ethics that come up, or more frequently online if necessary. The committee is also responsible for approving this deliverable.

We will have arranged for our first annual ethics review at the end of October with Dr Adam Henschke from the University of Twente in the Netherlands. Dr Henschke has a broad experience in research into ethical issues surrounding areas such as biotechnology and the military. His report will be taken into account in our ongoing work and will feed into the following ethics review.

The UTTER project will also establish a mechanism for providing 20% of the overall budget for small scale projects (financial support for third parties or FSTP), selected through two open calls. The objectives of these pilots should be to further the outcomes of the original EU Call (Call: HORIZON-CL4-2021-HUMAN-01 A HUMAN-CENTRED AND ETHICAL DEVELOPMENT OF DIGITAL AND INDUSTRIAL TECHNOLOGIES 2021 Topic: HORIZON-CL4-2021-HUMAN-01-13), i.e. to apply pretrained extended reality (XR) models to new applications and new sectors, to increase the usability and efficiency of these models and to test how they can enable new types of human-human and human-machine interaction. We will require that the projects selected follow the ethical guidelines of the UTTER project.

3 Protection of personal data

Our approach to processing data in UTTER includes a number of aspects specifically designed to increase the protection of personal information:

- Our approach to data management and the integration platform is based in part on the concepts of *privacy by design*² and by *legal protection by design*³.
- Any personally identifiable information (PII) that is not needed will be manually destroyed or anonymised.
- Some comments in meetings expressing personal views could be considered sensitive personal information (SPI) and will also be manually destroyed unless there is a compelling reason to retain it.

One of our key responsibilities is to ensure that the UTTER Data Management Plan (D1.1) complies with EU and national law. In particular, the EU General Data Protection Regulation (GDPR; <https://gdpr-info.eu>) became applicable in May 2018, replacing the EU Data Protection Directive, and we have ensured that UTTER is compliant with GDPR. Legal issues of specific relevance to UTTER include:

- Explicit descriptions of the data generated and collected in the project are provided in Section 1 of the Data Management Plan (UTTER deliverable D1.1).

² https://en.wikipedia.org/wiki/Privacy_by_design

³ See e.g. Hildebrandt (2017)

- Data Protection Impact Assessment – this has been applied for through the University of Edinburgh Data Protection Officer and has been approved.
- The project follows a process of Data Protection by Default and by Design, as mandated by article 25 of the GDPR:

implement appropriate technical and organisational measures, such as pseudonymisation, which are designed to implement data-protection principles, such as data minimisation, in an effective manner and to integrate the necessary safeguards into the processing [...] by default, only personal data which are necessary for each specific purpose of the processing are processed

- The project will follow an approach of *privacy by design* (and *legal protection by design*)⁴. This means that we shall endeavour to be aware of potential legal and ethical issues that may arise from the use of the meeting platform, and have considered these in the platform’s design (although actual implementation of such features is beyond the current project).
- The consortium will specify how we identify where personal data is involved and how such personal data is protected. As the UTTER architecture and information flow is established, security and privacy issues will be taken into consideration. Some consortium partners dealing with data, including provision, use, processing and storing, have looked into data protection regulations for their organisation and country.
- Some consortium partners have sought advice from their respective local data protection authorities (UK, France, Portugal, Netherlands). It is understood that the consortium as a whole are joint data controllers in this project.
- Security procedures have been established for each partner dealing with data. For internal project data, we will use a secure private GitLab instance hosted by the University of Amsterdam. This repository is being used to host data which contains PII and therefore cannot be released.
- The consortium is striving for transparency to make the purpose of including content containing personal data in the UTTER project clear. For instance with meeting recordings, we explicitly informed the meeting participants and obtained consent to record by getting them to sign an online form.
- During dissemination events, when targeting potential new users, the platforms will run with synthetic data containing no personal data content.

4 Experimental data

In order to examine the usefulness of our meeting assistant prototypes, we will use human participants with relevant language expertise to evaluate our models. We will follow the latest ethical processes to conduct this research. Following the ethics work package leader’s research ethics procedures (UEDIN), detailed here <http://web.inf.ed.ac.uk/infweb/research/ethics-and-integrity>, we

⁴ Danezis et al., “Privacy and Data Protection by Design – from policy to engineering”, ENISA (European Union Agency for Network and Information Security) report, December, 2014. <https://www.enisa.europa.eu/publications/privacy-and-data-protection-by-design>, see also Danezis et al. (2015).

make sure that participants are informed of what the experiments involve. We applied for ethics approval and have been approved.

In particular participants will be informed of the following:

- This is a research study and results will likely be used in publications.
- What the purpose of the research is.
- Employees must be allowed to volunteer freely.
- Personal information, such as names, will in no way be connected with the generated experimental data.
- The conditions under which their data will be stored.
- What they will be asked to do.
- If the participants have any ethical concerns and/or complaints they will be informed who they can contact. This could be the ethics person in Naver Labs Europe/Unbabel, or it can be the University of Edinburgh Informatics Ethics Committee, or both.

5 Broader ethical concerns

A number of broader ethical issues relating to UTTER have been discussed at the Ethics committee meetings. Although many of the issues discussed are broad societal issues, we have attempted to foresee potential ethical concerns.

1. Job security and automation.

The UTTER platform certainly has the potential to reduce the number of jobs directly concerned with, for example, customer service and translation. Ultimately our technologies are not meant to replace humans but empower humans in interactions with other humans. For instance the customer service use case doesn't replace service agents, but enables them to carry out more conversations more efficiently. Our tools also make meeting data more accessible for analysis and insight - ultimately leading to better understanding of the customer and better service. More broadly, the UTTER platform, has the potential to open new business areas. To address this issue more rigorously, requires research beyond that carried out by UTTER.

2. Transparency of automatic processing systems.

This is an extremely important issue, however in our case the original recording and language data is still provided to the user so any errors can be tracked down and addressed. We are also researching uncertainty quantification and explainability to improve transparency of the decisions take by the models themselves.

3. Dual use (military and defence).

As an EU project we use the EU definition of dual use, where the duality is purely civilian vs military. (Broader definitions of dual use refer to any malicious unintended reuse⁵)

⁵ see <https://arxiv.org/abs/2304.08315>

The project has no direct plans to have any military use case. However, as the tools are free/open-source, these could potentially be used by the military. It is clear that a good “meeting assistant” can also speed up and scale up military surveillance of meetings. We do not include a full analysis of this at this point, but we are committed to enumerating possible dual use risks of our research in publications and subsequent reviews, and where possible, laying out steps for preventative measures.

We believe the advantages of releasing our software as free/open-source for the benefit of researchers and businesses across Europe outweigh the risks.

6 Conclusion

The aim of this report is to provide a presentation of the key ethical questions and challenges that the UTTER project must address, our process for handling these issues, and the project’s responses to them.

References

George Danezis, Josep Domingo-Ferrer, Marit Hansen, Jaap-Henk Hoepman, Daniel Le Metayer, Rodica Tirtea, and Stefan Schiffner. Privacy and data protection by design — from policy to engineering. Technical report, 2015. arXiv preprint arXiv:1501.03726.

Mireille Hildebrandt. Saved by design? the case of legal protection by design. *Nanoethics*, 11(3): 307–311, 2017.

ENDPAGE

UTTER

HORIZON-CL4-2021-HUMAN-01 101070631

D9.1 UTTER First Ethics Review