The University Of Sheffield.

UKRI UK Research and Innovation

UKRI Centre for Doctoral Training in Speech and Language Technologies and their Applications

SLT
Centre for Doctoral Training

What are Speech and Language Technologies?

Since the inception of Artificial Intelligence (AI) in the 1950s, a central vision of the field has been of machines that can understand spoken and written human language.

This capability would enable natural interaction between people and computers, translation between all human languages, and tools to analyse and answer questions about vast archives of text and speech.

Speech and Language Technologies (SLTs) are a range of approaches which allow computer programs or electronic devices to analyse, produce, modify or respond to human texts and speech.

SLTs are underpinned by a number of fundamental research fields including natural language processing, computational linguistics, mathematics, machine learning, physics, psychology, computer science, acoustics, and speech processing.

In the last two decades the societal shift to digital media, coupled with spectacular advances in computer storage, processing power and software capabilities, especially in machine learning and SLTs, has meant this vision is no longer science fiction but is turning into reality.

SLTs are now established as core scientific/engineering disciplines within AI and have grown into a world-wide multi-billion dollar industry, with massive application potential (global revenues for the industry were valued at $33bn in 2015 and predicted to rise to $80bn by 2020, an annual average growth rate over 19%).

There is huge demand for scientists with advanced training in SLT from both large corporations and SMEs; most hire only at PhD level and are in fierce, international competition to attract talent: supply is falling far short of demand.

A PhD. With a difference.

Our Centre for Doctoral Training (CDT) addresses the SLT skills gap – the first ever dedicated CDT in this area.

And it goes far beyond a 'standard' PhD.

You will receive training in the core technical and scientific SLT skills that industry and the UK economy demand. You will be embedded in a vibrant research centre that provides training in engineering skills, leadership, management, entrepreneurship, ethics, and responsibility to society.

Our unique PhD with integrated Postgraduate Diploma (PGDip) in SLT Leadership is a bespoke programme that will provide you with the necessary skills for academic and industrial leadership in the field.

Furthermore, historically the individuals choosing to study Computer Science, and AI in particular, at the advanced level have been predominantly male and from majority ethnic groups. The SLT academic community is committed to attracting more women and minorities into AI and SLT research (some of these initiatives are led by our own CDT's academics). Building on these initiatives we will strive to attract a diverse set of students into the CDT.

The Centre

Our Centre for Doctoral Training is supported by UK Research and Innovation (UKRI), the UK's national funding organisation for research.

It is hosted by the University of Sheffield's Department of Computer Science which has an international reputation for the quality of its research and teaching. In the 2014 Research Excellence Framework (REF), 92% of our research work was rated world leading (4*) or internationally excellent (3*) in terms of its originality, significance and rigour. We have a REF grade point average (GPA) of 3.39, ranking us 5th out of 89 computer science departments in the UK.

We are training the next generation of highly skilled researchers in SLTs to gain the skills required to successfully enter the three main areas of need for SLT experts in the UK:
1. Academia
2. Mid-to-large scale SLT R&D teams
3. Start-ups

You will undertake high quality research that is relevant to the needs of industry alongside a comprehensive training programme addressing core and professional skills – gaining a Postgraduate Diploma (PGDip) in SLT Leadership integrated with your PhD.

The CDT offers a unique combination of industry-driven projects covering all areas of natural language processing (NLP) and speech processing research, including natural language analysis and generation, information retrieval, text mining (including sentiment analysis), question answering, machine translation, speech and speaker recognition, diarisation, machine hearing, novel methods of interaction and dialogue, and detection and analysis of paralinguistics.

Every PhD project is underpinned by a real-world application, ideally directly supported by an industry project partner.

As a student in the CDT, you will have a far superior level of support and training to that of a 'conventional' PhD studentship. You will receive 4 years of enhanced funding (including fees and a tax-free bursary larger than the UKRI standard), an increased personal research budget, specific training on core expertise, industry-led teaching, CDT-wide cross-disciplinary mini-projects, and much more.

To promote cohort cohesion we apply a uniform funding model for all students (same stipend, same training budget) regardless of the sponsorship approach adopted by the industrial partner with whom you work.

The CDT will train 60 university-based PhD students, split into five annual cohort intakes, over a total period of 8 years. Our first cohort of students started in September 2019.

The unique nature of the CDT means that you will gain both excellent academic research skills and highly valuable experience of addressing real-world research problems - making you a valuable asset to the industry.

Why Now?

Artificial Intelligence (AI) is a disruptive technology and is underpinning significant change in the global economy – a trend which will continue for the foreseeable future.

Recent reports by PwC (*The macroeconomic impact of artificial intelligence* and *The economic impact of artificial intelligence on the UK economy*) predict that by 2030 AI could contribute a 14% increase to global GDP, equivalent to $15.7 trillion. They see 70% of that additional value coming from North America and China. In the UK they predict GDP will be up to 10.3% higher in 2030 as a result of AI – the equivalent of an additional £232bn, leading Northern European economies just slightly.

To advance its position in the global AI race, the UK must invest in the primary resource needed to develop new AI technologies and enterprises: people with advanced training in AI.

Our CDT is playing a critical role in this national strategy to address the AI skills gap; SLT is particularly important for the following reasons:

1. SLT is a core area within AI because human language is absolutely central to all human social (and hence economic) activity and so the scope for application is virtually limitless (in the PwC report SLT and its applications figure centrally in their "layers" of AI).
2. SLT is an area where the UK is internationally leading, and it must invest to maintain its position, and where it benefits from the global dominance of the English language.
3. As the strong industrial support for this CDT shows, there is huge appetite within UK industry and the public sector for more skilled researchers in this area.

The Training Programme

Every student receives training in four domains:
1. core technical knowledge
2. SLT collaborative research practice
3. research and innovation professional skills
4. personal effectiveness and leadership

Training is delivered via several modes:
- individual study
- cohort-based study
- cohort-based project work
- external activities

Research streams are topically related themes that are designed to bring together researchers from all sub-areas of SLT for discussion and collaboration. Based on your project you will engage with at least 2 streams, and each stream is led by a co-investigator of the CDT with the objective of ensuring students gain a broad knowledge and coverage of the stream subject.

Informed by real world deployment of SLTs, the streams are:
- **I. SLT Frontiers – novel methods**: New machine learning and algorithmic approaches to solve SLT problems; novel interpretations of data; new methods of obtaining or characterising data.
- **II. Scalable SLTs**: SLT methods that allow adjustments of methods to different task complexities; transferring SLTs as component of larger tasks; design of scaleable SLTs.
- **III. Robust SLTs**: Resilience to SLT domain and data variability; methods to achieve uniform performance across diverse data; transfer learning and adaptability.
- **IV. Novel SLT applications**: Pushing the boundary on where SLT can be used; methods to obtain data for new applications; approaches for transferring of existing methods.
- **V. Interconnecting SLT with the world**: Approaches that connect advanced methods with real world problems; methods that allow cross-fertilisation between domains; transferable methods for SLT; methodology in approaching real world problems.

Student training is designed to be a journey that progressively expands the level of expertise in both research and professional, transferable skills. Training content is delivered in the form of credit-bearing, University-accredited modules.

The CDT demonstrates enhanced training beyond pure SLT research, via the Postgraduate Diploma (PGDip) in SLT Leadership, ensuring students obtain leadership, entrepreneurial, and software skills.

Training is tailored to each individual student's needs with individual module selection informed by the University of Sheffield Doctoral Development Programme and a training needs analysis (TNA) carried out with each new student during the first few weeks of their programme.

The Student Journey

Training Year 1 – Skill Foundation

SLT is exceptional in the range of disciplines which it draws upon, from linguistics and phonetics through mathematics and computer science to signal processing and electrical engineering. The first year is therefore designed to ensure that the group of students enrolling from diverse academic backgrounds can develop into a well-integrated, self-supporting cohort. Students receive unconscious bias and Equality, Diversity and Inclusivity (ED&I) training to perpetuate an environment of fairness, equality, diversity and respect. After an induction phase, student PhD projects will be defined in discussion with the students, supervisors and industrial partners. Students will start work on their research topic while still receiving foundational training.

Training Year 2 – Scientific Foundation

The second year is devoted to developing advanced SLT research skills in practice, to perform the first foundational experiments and to formulate the plan for the PhD. The student will engage in further cohort and external activities as well as receive further training in all training domains and modes.

Training Year 3 – Research

This is expected to be the most productive research year. Activities will be similar to those conducted in Y2, however the students are expected to perform more leadership roles in cohort and team work, e.g., by supervising mini-projects, by stepping into planning roles in the cohort-wide activities, or by mentoring of peers. Internships are likely to happen in Y2 or Y3.

Training Year 4 – Consolidation, Presentation and Dissemination

In the final year the emphasis will be on thesis completion and on ensuring impact through presentation or realisation in practical settings. The year will see the completion of the PGDip, followed by submission and assessment of the PhD thesis.

Cohort-wide training activities

Our CDT's students are trained using a cohort-based approach since:

1. the software infrastructure, tools and methods for SLT are highly complex and creating them is nearly always a collaborative endeavour – a cohort offers an ideal setting to gain experience of such collaborative working.
2. PhD topics tend to be narrow and focused on specifics and do not include the broad overview needed in students' later careers – through cohort training we can expose students to a range of different SLT topics.
3. peer learning within and across cohorts is a highly effective way to hand over tools and to teach methodology.
4. a multi-year cohort programme allows significant and sustained progression in larger (i.e. multi-student) SLT projects, resulting in better research outcomes and more impact in partnering companies.
5. an extended cohort-based training programme with strong group work and peer tutoring elements allows students with non-standard backgrounds be admitted, helping to promote diversity in SLT.

The CDT features a number of novel training elements to ensure you gain wide ranging professional skills and experience in activities relating to industry and academic life. It is of particular importance for the fast-paced world of AI and SLT to understand how to kickstart business, or work in effective academic or R&D teams.

The programme requires the completion of 120 credits of modules over the four years. In the first year, you will study 75 credits: two 15-credit core modules, plus three 15-credit optional modules. The three optional modules can be drawn from six of our host department's masters-level modules: Text Processing; Speech Processing; Machine Learning And Adaptive Intelligence; Scalable Machine Learning; Speech Technology; Natural Language Processing. In your second, third and fourth years you will study one 15-credit core module per year.

The training elements in each core module appear, in different forms, throughout your time in the CDT. Apart from the Personal Development Project (PDP) each element is designed allow cohorts to engage in different roles for the same activity.

Mini project (Year 1)

The Mini Project will lead you through the phases of creating an SLT research prototype system, from scoping through implementation, evaluation and dissemination. Projects will be chosen to span both speech and language sub-areas. Example topics are speech to speech translation, spoken question answering or call centre analytics. You will work as a cohort with each member undertaking different roles at different times. The module leader will meet students on a regular basis to review progress and provide suggestions. More senior cohorts (or existing PhD students in the first instance) will offer more detailed supervision/ technical assistance where appropriate.

The mini projects are chosen to provide experience of design and coding problems in Research Software Engineering (RSE) and to require the use of specific RSE software and hardware resources likely to be needed by students in their PhD projects.

SLT Hub (Years 2, 3, and 4)

This is a student-led SLT consultancy-style service to local and UK companies and other organisations. Short projects will be planned and delivered by small student teams formed to bring together the necessary expertise.

This activity is modelled on a unique feature of our host department: our undergraduate students run their own software company, Genesys, which accepts real contracts for real clients.

Hub projects are relatively small in size, and cover specific advice and small scale feasibility study work in the domain of speech and language technologies. You will take different roles engaging with partners and the wider SLT industry in the UK and overseas. All work is conducted in teams and you will be involved in the delivery of work, in the engagement with industry and in the planning and presentation work.

SLT Challenges (Years 2 and 3)

The SLT research community has pioneered the use of international shared task challenges to enable comparable and reproducible research. Teams work on common data sets and task definitions to compare the performance of algorithms under controlled conditions. Such challenges provide an excellent opportunity to teach core technical skills as well as scientific methodology and rigour.

As part of a CDT challenge team, you will organise and execute a complete challenge once a year, from data collection and preparation to distributions of baseline algorithms, scoring tools and assessment of outcomes, and publication in appropriate venues.

**Personal Development Project (Years 1, 2, 3, and 4)**

In year 1, you will choose a small project that is close to your interest, but distinct from your PhD project. This may be an outreach project, e.g. teaching AI in schools, an open source software project, or a data collection project in an under-resourced domain. Over the 4 years, you will plan the project, find partners to work with, and even obtain small amounts of funding if necessary. This is an opportunity for you to think longterm and learn in practice about the realisation of SLT and related projects.

SLT Journal and Seminar Club (Years 1, 2, 3, and 4)

Keeping up with the literature is a vital component of a doctorate and an important aspect is expert discussion of such material. The club will mix student and external talks with reading of relevant literature relating to the research streams. The club will include peer presentations – presentations by one student of the work of another – to train you in transferring of knowledge and gain a deeper understanding of work by others. Journal clubs will be theme specific and students are required to read on at least two themes.

SLT Meta MOOC (Years 1, 2, 3, and 4)

A lot of useful SLT learning material is available on the Web in the form of Massive Open Online Courses or MOOCs. The SLT Meta MOOC is a resource for you to learn about both general and highly specialised topics in SLT, created by you and your fellow students under the supervision of experts.
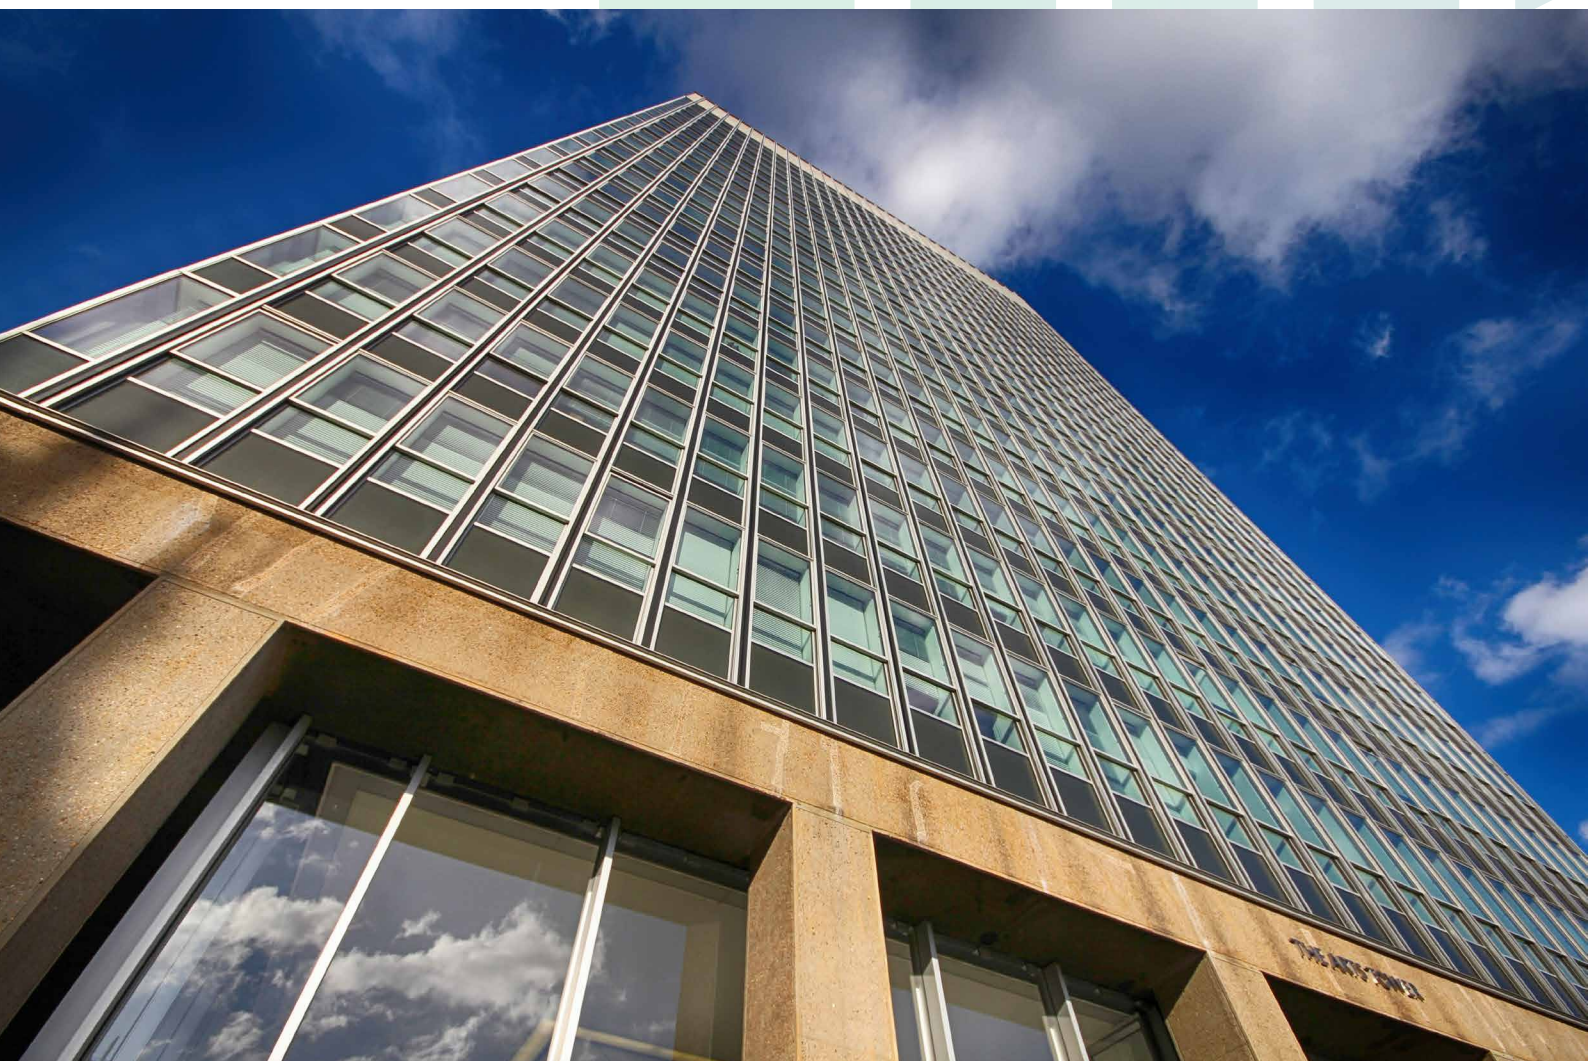
This is modelled on two aspects, the ISCA SCOOT resource (in which Sheffield plays a leading role), as well as the technical blogs on specific subjects, which are now becoming common in machine learning, vision and speech and language.

You will take up the role of topic scout, writer and/or reviewer and produce content referencing online lectures, papers and other online material. Material will be made available internally and, following quality assurance, released for public access. Theme leads and a module lead will support the selection of topics, balance and quality.

Supervision and Mentoring (Years 1, 2, 3, and 4)

An outside view is essential for productive work. Every year you will mentor a more junior student or co-supervise a Y1 mini-project or an SLT-related project of a departmental MSc or undergraduate student.

Annual CDT Training Workshop (Years 1, 2, 3, and 4)

Each year in Month 6 the CDT will run a workshop that will include training on topics including: responsible research and innovation; commercialisation of research and entrepreneurship; research software engineering, in particular, tools and techniques of relevance to SLT. Workshop content will be delivered mostly by external partners of the CDT. All training will be offered across cohorts, and will include team elements.

Your research project

Your PhD research will be built on the foundations of an industry sponsor's technical and commercial needs. You will have the opportunity to work closely with them; many research students go on to join their sponsors in key technical roles.

The CDT team will work with a potential sponsor to identify medium term research projects that will benefit their organisation's activities, but which also have academic depth suitable for doctorate study. You will help define the project, be provided with experimental data and/or access to facilities, receive supervisory input from your sponsor via periodic progress meetings, and have the opportunity to undertake an internship.

Both you and the sponsor's staff will benefit from effective knowledge transfer through the CDT's activities. You will gain industry-specific skills and insights from your sponsor's team, and have the opportunity to see the results of your research integrated into their business processes.

You will also benefit from access to the wider cohort's industrial sponsors - their teams will be involved in your training providing you with unparalleled exposure to a wide range of application domains and company types.

How to apply?

Eligibility

The four year funding covers the annual university fees, a highly competitive (enhanced) tax-free stipend (£17,000 per year), and a training budget.

To be considered for one of the studentships, you will need to satisfy the UKRI / EPSRC funding eligibility criteria.

In particular, to be eligible for a full award (stipend and fees), you must have:
- settled status in the UK, meaning you have no restrictions on how long you can stay,
- have been 'ordinarily resident' in the UK for three years prior to the start of the studentship. This means you must have been normally residing in the UK (apart from temporary or occasional absences); and,
- not been residing in the UK wholly or mainly for the purpose of full-time education (this does not apply to UK or EU nationals).

Entry requirements

Applicants should have, or be expecting to obtain, a minimum of a 2:1 undergraduate degree or a masters degree (ideally distinction) in a relevant discipline.

Suitable backgrounds are (but not limited to):
- Computer Science
- Engineering
- Linguistics
- Mathematics
- Physics
- Psychology

Regardless of background, you must be able to demonstrate mathematical aptitude (minimally to A-Level standard or equivalent) and experience of programming. We will also consider applicants with a professional background, so long as you are able to provide evidence of demonstrable academic skills as well as practical experience.

If English is not your first language, you will need to meet our English Language Requirements. We ask for IELTS 7.5 overall, with no less than 7.0 in each component.

Apply

Applications to the CDT are made through the University of Sheffield's Postgraduate Online Application Portal: https://www.sheffield.ac.uk/postgradapplication/

You will need to provide proof of your previous degree(s), your CV, two academic references, and a supporting statement describing your interest in the CDT and, if possible, the research area that interests you.

> To find out more about joining the CDT or if you have queries regarding the process or your eligibility, please email sltcdt-enquiries@sheffield.ac.uk

The Team

Directors

Thomas Hain
Head of Speech and Hearing Research

World leader in speech recognition, heads the VoiceBase Centre for Speech and Language Technology, and is a leader in the speech community.
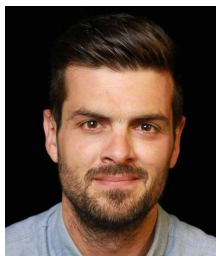
Rob Gaizauskas
Natural Language Processing Research

Internationally known for his research on information extraction and text mining, temporal information processing, question answering and summarisation.

Theme Leads

Nikos Aletras
Theme: SLT Frontiers - novel methods

Expert in topic detection and interpretation in large volumes of text data.

Jon Barker
Theme: Robust SLTs

World leader in research into speech communication in natural environments.

Loïc Barrault
Theme: Scalable SLTs

Expert in statistical and neural machine translation including linguistics aspects (factored neural machine translation) and considering multiple modalities (multimodal neural machine translation).

Stefan Goetze
Theme: Interconnecting SLT with the world

Expert in sound processing and enhancement; assistive technologies; human–machine interaction; detection and classification of acoustic events; and automatic speech recognition.

Heidi Christensen
Theme: Novel SLT applications

Expert in a wide variety of clinical application of speech technology, spoken language processing and binaural machine listening.

Supervisors

- **Guy Brown** is world renowned for his work on Computational Auditory Scene Analysis and his research into noise robustness of speech technology and hearing impairment.

- **Paul Clough** is an expert in information retrieval, specifically multilingual and geographical search, image retrieval, text re-use and plagiarism detection, and evaluation of search systems.

- **Yoshi Gotoh** is an expert in interpreting speech in the context of audio visual processing, in particular, applications of speech recognition in video analysis and video information retrieval.

- **Mark Hepple** is an NLP expert in many areas with specialisation in formal grammar and parsing, information extraction, temporal information processing, and robust dialogue processing.

- **Chengua Lin** is an expert in the development of algorithms and models for natural language generation, sentiment analysis, text/opinion summarisation, and intention recognition for cybersecurity.

- **Diana Maynard** is lead computational linguist directing the development of the GATE framework, amd expert in terminology extraction and ontology development.

- **Roger Moore** was president of the International Speech Communication Association (ISCA) and is a highly regarded international speaker with a keen interest in bringing SLTs to robotics.

- **Anton Ragni** is an expert in kernel methods for speech processing, discriminative modelling for speech processing, and generative modelling for speech processing.

- **Carolina Scarton** is an expert in text adaptation, machine translation, online misinformation detection and verification, evaluation of NLP task outputs, NLP applied to healthcare and robotics, and dialog systems.

- **Lucia Specia** is an expert in machine translation (MT) and text adaptation; currently she holds an ERC fellowship to investigate the use of multi-modal, contextual information to improve MT.

- **Mark Stevenson** is internationally recognised for his work in word sense disambiguation, information extraction and retrieval, plagiarism detection, lexicon adaptation, exploratory search.

- **Aline Villavicencio** is an expert in computational models for languages, using a combination of machine learning and linguistic/psychological theories to tackle issues of multilinguality, including for low-resourced languages in applications such as machine translation and text simplification.

In addition, affiliate supervisors are also drawn from various associated research groups both from within the Department of Computer Science as well as other University of Sheffield departments and research institutes.

UKRI CDT in Speech and Language Technologies
Partners and Affiliates

The University Of Sheffield.

UKRI — UK Research and Innovation

Apple	SOLVAY
Google	Tech Nation
Amazon Research	Therapy Box
VoiceBase	Gweek
Microsoft Research	Signal A.I.
Nuance	Textio
MapR	NHS Digital
ZOO Digital	OCLC
Emotech	TribePad
Scribetech	SoapBox Labs
Factmata	Jam Creative Studios
Ieso Digital Health Ltd	Netcall
ITSLanguage	Sheffield Digital
3M	King's College London
Ontotext	BTS
Recordsure	Kollïder

Want more information?

🌐 For more information, please visit **slt-cdt.ac.uk**

@ Join our email list - email your details to **sltcdt-enquiries@sheffield.ac.uk**

🐦 Follow us **@sltcdt**

SLT Centre for Doctoral Training