

Lecture & Date	Topic	Presented papers
Introductory classes (Week 1-2)		
Lecture 1 01.09.2026	Introduction to the course	
Lecture 2 03.09.2026	Multilingual transformer models	
Lecture 3 08.09.2026	Introduction to Machine translation	
Lecture 4 10.09.2026	Introduction to Speech processing	
Review topics & papers		
Week 3		
Lecture 5 15.09.2026	Tokenization Lead discussion: Paper 1: Stephanie-Ye Xia, William Bourgeois Paper 2: Ruth Crasto, Jacob-Junqi Tian	Language Model Tokenizers Introduce Unfairness Between Languages Do All Languages Cost the Same? Tokenization in the Era of Commercial Language Models (Presenter: Soraya Ghassemlou & Cher Li) Zero-Shot Tokenizer Transfer (Presenter: Victor Akinode) Further reading: Julie Kallini's slides
Lecture 6 17.09.2026	Cross-lingual Transfer Lead discussion: Paper 1: Luis Lara Paper 2: Tong Wu & Ezra Huang	Beto, Bentz, Becas: The Surprising Cross-Lingual Effectiveness of BERT Expanding Pretrained Models to Thousands More Languages via Lexicon-based Adaptation (Damaris Stephanie Ndjebayi & York Hay Ng) XLM-V: Overcoming the Vocabulary Bottleneck in Multilingual Masked Language Models (Shaani Bellemare & Hao Yuan Lu) Additional reading: Extending Multilingual BERT to Low-Resource Languages Revisiting Machine Translation for Cross-lingual Classification

Week 4		
Lecture 7 22.09.2026	Sentence embedding Lead discussion: Paper 1: York Hay Ng Paper 2: Michael Yu	Text Embeddings by Weakly-Supervised Contrastive Pre-training M3-Embedding: Multi-Linguality, Multi-Functionality, Multi-Granularity Text Embeddings Through Self-Knowledge Distillation (Ezra Huang & John Dellas) Improving Text Embeddings with Large Language Models (Stephanie-Ye Xia & Ruth Crasto) Additional reading: MMTEB: Massive Multilingual Text Embedding Benchmark Language-agnostic BERT Sentence Embedding
Lecture 8 24.09.2026	Machine Translation Lead discussion: Paper 1: Sumukh Pate & Cher Li Paper 2: Jiaduo Xing & Shaani Bellemare	Scaling neural machine translation to 200 languages Multilingual Machine Translation with Open Large Language Models at Practical Scale: An Empirical Study (Marie Maltais & Luis Laramata) NiuTrans.LMT: Toward Inclusive and Scalable Multilingual Machine Translation with LLMs (Henok Ademtew & Jacob-Junqi Tian) Additional reading: X-ALMA: Plug & Play Modules and Adaptive Rejection for Quality Translation at Scale
Week 5		
Lecture 9 29.09.2026	Machine Translation Evaluation Lead discussion: Paper 1: Yilin An, Winnie Xiong Paper 2: Henok Ademtew, Allenna Tang	COMET: A Neural Framework for MT Evaluation MetricX-24: The Google Submission to the WMT 2024 Metrics Shared Task (Jiaduo Xing) SSA-COMET: Do LLMs Outperform Learned Metrics in Evaluating MT for Under-Resourced African Languages? (Marzia Nouri) Additional reading: Large Language Models Are State-of-the-Art Evaluators of Translation Quality
Lecture 10 01.10.2026	Ethical data collection & participatory research Lead discussion: Paper 1: Soraya Ghassemlou, Marie Maltais Paper 2: David Guzmán Ramírez, Hao Yuan Lu	Building Better: Avoiding Pitfalls in Developing Language Resources when Data is Scarce Participatory Research for Low-resourced Machine Translation: A Case Study in African Languages (William Bourgeois & Michael Yu) SEACrowd: A Multilingual Multimodal Data Hub and Benchmark Suite for Southeast Asian Languages (Sumukh Patel) Additional reading: Global PIQA: Evaluating Commonsense Reasoning Across 100+ Languages and Cultures

Week 6		
Lecture 11 06.10.2026	Guest lecture	Prof. Skyler Wang (McGill & Handshake AI): Ethical data collection and Scaling NLP to several low-resource languages
Lecture 12 08.10.2026	LID & Multilingual pre-training data Lead discussion: Paper 1: Marzia Nouri, Ankur Sikarwar Paper 2: Mohamed Elsamadouny	What Language is This? Ask Your Tokenizer Quality at a Glance: An Audit of Web-Crawled Multilingual Datasets (Yilin An & Winnie Xiong) FineWeb2: One Pipeline to Scale Them All — Adapting Pre-Training Data Processing to Every Language (David Guzmán Ramírez & Shawn Hu) Additional Reading: GlotLID: Language Identification for Low-Resource Languages MADLAD-400: A Multilingual And Document-Level Large Audited Dataset
Week 7		
13.10.2026	Reading week	
Lecture 13 15.10.2026	LLM architecture Lead discussion: Paper 1: Victor Akinode, Damaris Stephanie Ndjebayi Paper 2: John Dellas, Shawn Hu	EuroLLM-9B: Technical Report OLMo: Accelerating the Science of Language Models (Ankur Sikarwar & Tong Wu) Apertus: Democratizing Open and Compliant LLMs for Global Language Environments (Mohamed Elsamadouny & Allenna Tang) Additional Reading: Olmo 3 TinyAya
Week 8		
Lecture 14 20.10.2026	LLM adaptation Lead discussion:	SEA-LION: Southeast Asian Languages in One Network AfriqueLLM: How Data Mixing and Model Architecture Impact Continued Pre-training for African Languages Instructing Large Language Models for Low-Resource Languages: A Systematic Study for Basque
Lecture 15 22.10.2026	Guest lecture	Dr. Julia Kreutzer (Cohere Labs): Multilingual reasoning
Week 9		
Lecture 16 27.10.2026	LLM reasoning Lead discussion: Stephanie-Ye Xia	Language Models are Multilingual Chain-of-Thought Reasoners Question Translation Training for Better Multilingual Reasoning (Ankur Sikarwar & Tong Wu)

		Long Chain-of-Thought Reasoning Across Languages Winnie Xiong & Yilin An Additional Reading: PolyMath: Evaluating Mathematical Reasoning in Multilingual Contexts
Lecture 17 29.10.2026	Project presentations (5 minutes each)	All students presents their project (in group of one or two) - in person participation only
Week 10		
Lecture 18 03.11.2026	LLM-as-a-Judge Lead discussion:	Lower-Resource, Higher Scores: Language Bias in LLM Evaluators Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena (Michael Yu & York Ng) Are Large Language Model-based Evaluators the Solution to Scaling Up Multilingual Evaluation? (Jacob-Junqi Tian & Jiaduo Xing) Additional Reading: Challenges and Recommendations for LLMs-as-a-Judge in Multilingual Settings and Low-Resource Languages
Lecture 19 05.11.2026	LLM safety Lead discussion:	JailbreakBench: An Open Robustness Benchmark for Jailbreaking Large Language Models Multilingual Jailbreak Challenges in Large Language Models (Victor Akinode & Henok Ademtew) MPO: Multilingual Safety Alignment via Reward Gap Optimization (Allenna Tang & Marzia Nouri) Additional Reading: The State of Multilingual LLM Safety Research: From Measuring The Language Gap To Mitigating It All Languages Matter: On the Multilingual Safety of LLMs
Advanced Topics		
Week 11		
Lecture 20 10.11.2026	Multimodal Multicultural LLMs	Multilingual Diversity Improves Vision-Language Representations Improved Baselines with Visual Instruction Tuning (Damaris Stephanie Ndjebayi & William Bourgeois) Jagle: Building a Large-Scale Japanese Multimodal Post-Training Dataset for Vision-Language Models (Stephanie-Ye Xia & Ruth Crasto) Additional Reading:

		BLEnD-Vis: Benchmarking Multimodal Cultural Understanding in Vision Language Models X-LLaVA: Optimizing Bilingual Large Vision-Language Alignment
Lecture 21 12.11.2026	Guest lecture	Dr. Min Ma (Google DeepMind): Speech processing
Week 12		
Lecture 22 17.11.2026	Multilingual Speech Representations	Towards Robust Speech Representation Learning for Thousands of Languages wav2vec 2.0: A Framework for Self-Supervised Learning of Speech Representations (Luis Lara & Marie Maltais) Scaling Speech Technology to 1.000+ Languages (Soraya Ghassemilou & Cher Li) Additional reading: XLSR: Unsupervised Cross-lingual Representation Learning for Speech Recognition HuBERT: Self-Supervised Speech Representation Learning by Masked Prediction of Hidden Units
Lecture 24 19.11.2026	Speech-to-Speech Translation	SeamlessM4T: Massively Multilingual & Multimodal Machine Translation UnitY: Two-pass Direct Speech-to-speech Translation with Discrete Units UNISS: Unified Expressive Speech-to-Speech Translation with Your Voice Additional reading: Direct speech-to-speech translation with a sequence-to-sequence model
Lecture 23 19.11.2026	AudioLLMs	SALMONN: Towards Generic Hearing Abilities for Large Language Models AudioGPT: Understanding and Generating Speech, Music, Sound, and Talking Head Omni-R1: Do You Really Need Audio to Fine-Tune Your Audio LLM? Additional Reading: Do Audio LLMs Really LISTEN, or Just Transcribe? Measuring Lexical vs. Acoustic Emotion Cues Reliance
Week 13: Class Project Presentations		
Lecture 24 24.11.2026	Final presentations (12 minutes talk + 3 minutes Q/A)	

Lecture 25 26.11.2026	Final presentations (12 minutes talk + 3 minutes Q/A)	
Week 14: Class Project Presentations		
Lecture 26 01.12.2026	Final presentations (12 minutes talk + 3 minutes Q/A)	