PROJECT ANNEX

Tatjana Chernenko | Applied AI & Data Scientist | AI Architect | Enterprise AI Advisor

Executive Profile

Principal-level AI & Data Scientist, AI Architect, and Enterprise AI Advisor with 20 years of overall professional experience, including more than a decade of hands-on work across Speech AI, GenAI, LLMs, Agentic AI, AI Speech Technologies, RAG, NLP, multilingual language technologies, data strategy, applied research, and production-oriented implementation.

Built credibility across the full lifecycle of enterprise AI systems: problem framing, applied research, prototyping, evaluation and benchmarking, data and pipeline architecture, production-oriented engineering, stakeholder consulting, and organisational AI enablement.

Particularly strong at turning ambiguous or emerging AI opportunities into technically credible and operationally realistic enterprise directions. Retains hands-on depth in Speech AI, speech translation, specialized terminology handling, multilingual and conversational systems, retrieval and evaluation, RAG, knowledge-graph-enhanced architectures, hybrid AI architectures, agentic AI workflows, data-centric AI design, and future-oriented intelligent systems including robotics-related scenarios.

Combines scientific depth, architecture thinking, and enterprise execution realism: able to shape applied research directions, design AI-ready data and evaluation foundations, and help organizations make better AI decisions under platform, legal, process, and delivery constraints.

Note: Across multiple initiatives, acted as a de facto owner of research direction, evaluation logic, data and AI architecture, and cross-functional enterprise execution.

Core Strengths

Primary focus

- Applied AI Research
- Speech AI / ASR / TTS
- Multilingual AI / NLP
- Evaluation, Benchmarking & Model Quality
- Domain Adaption
- Retrieval, RAG
- Agentic AI
- Knowledge Graphs
- Quality Imporvement
- Complex Data & Evaluation architectures & pipelines

Strong secondary focus

- AI system design for language and speech workflows
- Data-centric AI architectures for enterprise use cases
- Retrieval- and agent-based systems

Additional value

- Research leadership
- Enterprise AI advisory, stakeholder alignment, AI direction-setting under ambiguity

- Enterprise applicability under constraints
- Workshop facilitation
- Cross-functional translation
- Stakeholder enablement
- Patents, publications, mentoring, AI capability building, and international research collaboration

Selected Impact Highlights

- **Enterprise Speech AI (Applied Research & Evaluation, Architecture & Leadership)**: Led the expansion of SAP's internal AI scope from traditional language technologies to a production-grade Speech AI capability, defining the applied research direction, models and methods selection, evaluation pipelines, end-to-end AI architecture and SAP-specific adaptation path across speech-to-text, subtitling, translated transcription, voice-over, multilingual speech processing, and speech-to-speech translation. Owned AI direction as the sole AI expert, combining technical direction with hands-on quality framing and evaluation-driven optimisation and production enablement, delivering internal multilingual speech workflows with ~10% MIM improvement over external solutions.
- **Speech AI Evaluation & Benchmarking Architecture:** Built the enterprise Speech AI (STT, TTS, Speech Translation) evaluation and benchmarking framework from scratch as the primary AI owner, creating a reusable benchmark asset spanning ~90 datasets, ~12 languages, and 100+ evaluation dimensions across meetings, subtitling, media, and speech-interface scenarios. Defined the methodology, benchmark structure, and quality criteria; built evaluation logic for comparative model assessment and quality diagnosis; and led evaluation-driven optimisation and stabilisation, achieving up to 20% Word Error Rate improvement. Co-uthored a patent on hallucination and skipped context reduction and alignment for speech-to-text pipelines.
- **TTS Optimisation & Patentable Innovation (Japanese)**: Owned a Japan-focused applied Speech AI direction to address enterprise-critical issues: unnatural rendering of technical terminology, anglicisms, and abbreviations in Japanese in settings where extensive fine-tuning was impractical due to data sensitivity, compliance overhead, and cost. Personally drove the problem framing, method qualification, and experimental direction for terminology- and pronunciation-focused improvement under real enterprise constraints, validating lower-cost intervention paths where broad fine-tuning was impractical. Defined the solution path, mentored experimentation and evaluation in Japan, delivered up to 15% Word Error Rate improvement in early evaluation, and co-authored the **patent** *"Enhancing Pronunciation Accuracy and Naturalness in Japanese Text-to-Speech Systems without Extensive Fine-tuning"* (under review).
- **Speech Data Platform & Architecture:** Architected a **SAP-wide multilingual speech data product** from the ground up, defining the schema, metadata model, storage strategy, and **20+ stage** processing architecture required to transform fragmented, highly sensitive audio sources into governed enterprise AI assets. Designed for scale across **~40 languages** and **100+ metadata dimensions**, creating the data foundation for SAP-specific **STT, TTS, speech-to-speech, evaluation, subtitling, voice-over, robotics, and agentic AI** use cases.
- **Applied Research (Low-Resource Speech AI - UNISINOS)**: Owned the scientific framing and applied research direction for a strategic collaboration with UNISINOS focused on a critical enterprise speech-to-text AI challenge: recognising customer-specific terminology and specialised jargon in low-data environments without large domain-specific fine-tuning corpora. Defined the research methodology end-to-end - from literature review, method qualification, benchmark design, evaluation methodology, and experimental boudaries to execution oversight. Qualified promising methods from the literature; co-designed an open-vocabulary keyword spotting architecture using synthetic keyword audio and cross-attention for acoustic matching; and contributed to training strategy and hard-negative sampling design. The work yielded up to +25% MRR in English and +133% in Portuguese on selected

company-domain benchmarks and resulted in an accepted LREC 2026 paper on specialised vocabulary evaluation, and a publication-grade internal system paper covering the architecture and training approach that remained confidential due to product constraints.

- **AI Architecture & Feasibility Leadership (RAG, Query Reduction, Agentic AI):** Led the AI-side feasibility, research and architecture direction for a Query Reduction initiative with supplier GlobalizeMe, assessing how RAG, embeddings, and fine-tuning could reduce query volume and processing complexity in SAP language-service workflows involving ambiguous, high-volume enterprise data. Defined architecture, evaluation frameworks, and realistic implementation boundaries for advanced workflow automation
- **Principal AI Advisory to Leadership & Enterprise Decision-Making:** Acted as a principal-level AI advisor across Speech AI, LLMs, RAG, knowledge graphs, and agentic systems. Guided decision-making across 100+ AI use cases, contributed as the sole AI expert at a leadership offsite (Brazil), and improved decision quality by identifying low-feasibility directions across models, data, evaluation, and architecture.
- **Enterprise NLP Platform Transformation (SAP Translation Systems)**: Rebuilt the AI retrieval and evaluation backbone for SAP's multilingual translation platform (40+ languages, millions of translations) as the primary AI owner for retrieval relevance, evaluation architecture, and quality measurement. Drove the transition from keyword-based to semantic retrieval and implemented scalable benchmarking pipelines (36+ reports per cycle). Established a repeatable experimentation and measurement layer that improved proposal quality by up to 10% and enabled measurable, system-wide optimisation. Supported large-scale enterprise rollouts (80k+ text elements, 25k+ objects) and contributed to a platform whose customer case studies, including Siemens Healthineers and Daimler Truck, reported up to 93% reduction in manual translation effort and 99%+ proposal accuracy.
- **Frontier AI Research (Speech AI for Robotics)**: Contributed to international applied research on Speech AI for robotics and voice-enabled intelligent systems, supporting SAP-side work on multilingual speech input/output, model selection, evaluation design, and deployment-constrained optimisation for edge and offline scenarios. Worked across a high-profile collaboration involving SAP Japan, SAP Research & Innovation in Potsdam, TRUSCO Nakayama, AIMBO Robotics, and Nagoya University's Kawaguchi Laboratory, adding frontier-facing research breadth across robotics, embodied AI, warehouse automation, and S/4HANA-connected intelligent systems.
- **AI Governance & Execution (Enterprise AI Delivery):** Enabled enterprise AI execution in compliance-sensitive environments by translating legal, licensing, procurement, data-protection, and supplier constraints into workable technical paths for real AI delivery. De-risked initiatives involving highly sensitive enterprise audio data, defined practical governance processes for production use of open-source, commercial, and internal datasets, and established the Speech Data Circle to align standards and operational readiness for SAP-specific Speech AI. Established governance and data-readiness patterns that made sensitive enterprise Speech AI use cases technically executable under compliance, licensing, and supplier constraints. Led cross-team standardisation of data interchange across five core LX components, defining common formats and a shared processing library that improved interoperability, reduced communication overhead, and strengthened the architecture foundation for large-scale enterprise ML/NLP workflows.
- **AI Capability Building & Transformation Enablement**: Designed and co-delivered structured AI enablement across Germany and Brazil, helping technical, business, and domain-focused stakeholders build practical understanding of agentic AI, knowledge graphs, GenAI direction, and applied AI strategy. Strengthened organisational AI maturity by turning fragmented awareness into more aligned expectations, shared language, and realistic readiness for broader AI transformation.
- **Large-Scale Linguistic Data Intelligence & Patent-Backed NLP Innovation:** Conducted large-scale analysis of 12M+ multilingual SAP terminology pairs, identifying structural and metadata issues

blocking AI reuse. Co-authored multiple patents in semantic domain assignment, terminology disambiguation, and abbreviation resolution, strengthening enterprise NLP capabilities.

1. Selected Research & Innovation Output

- Co-author of **LREC 2026 paper**: "A Dataset for Evaluating ASR on Specialised Vocabulary" (Submission 1226), focused on evaluation methodology and benchmark design for specialised-vocabulary robustness in enterprise ASR.
- Inventor / co-inventor on **"Semantic Domain Assignment Referencing Governance Domains and Term Databases"** (US-12518105-B2, 6 January 2026) - patent covering automated assignment of semantic-domain labels to translation or terminology data by referencing governance-domain structures and term-database metadata, aimed at improving domain classification consistency and downstream reuse of enterprise language assets.
- Inventor / co-inventor on **"System and Method Performing Terminology Disambiguation"** (US-12386820-B2, 13 August 2025) - patent on terminology disambiguation using terminology-database lookups and metadata sequences to resolve ambiguous terms in context and communicate the intended meaning more reliably in enterprise language workflows.
- Inventor / co-inventor on **"Detection of Abbreviation and Mapping to Full Original Term"** (US-12067370-B2, 20 August 2024) - patent focused on detecting and resolving non-standard enterprise abbreviations - including atypical, concatenated, and camel-case forms - by mapping them to their full underlying terms to improve translation and language-processing quality where abbreviation handling would otherwise reduce precision or consistency.
- Inventor / co-inventor on **"Adaptive Fidelity Pipeline for Minimizing Hallucinations and Skipped Content in Speech-to-Text Systems"** (pending, 250089US01, Aug 18, 2025) - patent focused on a specialised preprocessing pipeline for speech-to-text systems that suppresses or removes input-side triggers associated with hallucinations and skipped content, improving transcription fidelity in challenging audio conditions.
- **Dialogue Management: Improving Task-Oriented Conversational Agents with Deep Learning**. Thesis in Computational Linguistics at Heidelberg University,Institute for Computational Linguistics supervised by Prof. Dr. Riezler and Dr. Vivi Nastase. Conversational AI-powered interfaces have become a top priority for many companies. Automated solutions save money, improve the quality of services and support the workers in their regular tasks. Technical software support is one of the industry fields that can benefit from usage of intelligent agents. AI systems with very specific knowledge in a narrow domain can automate regular tasks and support the workers by resolving technical issues. This thesis presents a development of such an AI-powered interface SmartSupport that combines strengths of the retrieval and neural models. SmartSupport can be used as an independent system or be integrated into a conversational agent as an underlying Question Answering System, improving its performance in a semantically restricted conversational space of technical software issues.
- Co-author of **Natural Language Generation from Structured Inputs for Image Description Generation (2018)** - Academic research project at Heidelberg University on structured-to-text generation for image description, aimed at verbalising object interactions, spatial relationships, action labels, and semantic-role structure. Built an attribute-to-text encoder–decoder pipeline with a feed-forward encoder over normalized structured inputs and an LSTM decoder for sequence generation, using MS COCO, V-COCO, and COCO-a to test how richer interaction-level representations affect caption quality under both automatic and human evaluation.
- Co-author of **Word Sense Induction for Web Search Result Clustering (CHERTOY system, 2017)** - Academic NLP research project at Heidelberg University built around the SemEval-2013 WSI benchmark. Designed and evaluated an unsupervised semantic clustering pipeline for ambiguity-aware organisation of web-search results, combining sense2vec-based lexical representations, vector-mixture

snippet embeddings, and MeanShift clustering, and ran 40 ablation-style experiments across preprocessing, language models, compositional setups, and clustering strategies; the final CHERTOY configuration materially outperformed the baseline on pairwise clustering metrics (RI 0.3732 → 0.6823; ARI 0.0284 → 0.1538; Jaccard 0.1431 → 0.6249).

2. Leadership Innovation and Global Research Co-Shaping

Period: 2025–2026 | **Role:** Research Direction Contributor / Innovation Steering Participant / Technical Advisor

Contributed to leadership-level innovation and research discussions across a portfolio of external, academic, and strategic AI initiatives, helping co-shape technical direction, research framing, and feasibility assessment across multiple topics and geographies.

Key Contributions

- Contributed to leadership-level innovation steering and research-direction discussions across 15+ completed initiatives.
- Acted as one of a group shaping technical direction, research framing, and feasibility assessment across a strategically relevant AI portfolio.
- Helped shape technical direction across a broad portfolio spanning speech and voice systems, machine translation and multilingual AI, knowledge-enhanced AI, retrieval-related topics, and structured-data-based generation.

Thematic List:

- **Speech and Voice Systems:**
- Customer jargon injection for Speech-to-Text research collaboration with University of Vale do Rio dos Sinos (UNISINOS)
- Indic Language Automatic Speech Recognition for SAP products in collaboration with IIIT Bangalore (building an ASR solution for HI-EN / indic languages for enterprise usage at SAP)
- Developing a Germany-based voice-assistance platform in collaboration with SPEAKER project funded by LX and German government
- Robotics-related applied research collaboration linked to Japan- and China-based initiatives and industry–academia exchange, including Trusco Nakayama, Aimbo Robotics, and Nagoya University's Kawaguchi Laboratory.
- **Machine Translation and multilingual AI:**
- Improve the quality of MT outputs in Simplified Chinese (joint research with SUFE on Machine Translation and Post-editing)
- Collaborative research with NICT on MT, specifically for lower-resource language pairs and text domains
- Learning from light-weight constraints to improve translation quality (PhD-work at University Heidelberg)
- Context MT: Improve MT for UI strings (PhD work at the Karlsruhe Institute of Technology)
- **Knowledge-enhanced and structured AI:**
- Generate texts from structured data (Master Thesis "Konzeption, prototypische Umsetzung und Evaluierung einer automatisierten Leiche Sprache Ausgabe aus WIKIDATA" in the Technische Hochschule Brandenburg)
- Evaluate application of NLP, AI and ML techniques to Query Management in collaboration with Pontifical Catholic University of Rio Grande do Sul (PUCRS)

Ownership

Acted as one of the small-number contributors helping shape research framing, technical direction, and feasibility assessment across a strategically relevant AI portfolio.

3. Principal AI Science and Applied Research

Enterprise Speech AI Research: Specialised Vocabulary, Jargon Injection, and Low-Data Adaptation (Academic and External Research Collaboration with UNISINOS)

Period: 2025–2026 | Role: Applied Research Lead / Research Direction Owner

Mission. Originated and led a strategic applied research direction, executed through an external collaboration with UNISINOS, to address a critical enterprise speech-to-text AI challenge: accurate recognition of customer-specific terminology and specialised jargon in low-data settings where large domain-specific audio corpora for fine-tuning were not realistically available.

Key Contributions

- Identified specialised vocabulary recognition and customer jargon injection as a priority enterprise Speech AI problem and elevated it from a product pain point into a structured applied research initiative with clear scientific and technical direction.
- Led the pre-research phase from the company side, including scientific literature review, method qualification, assessment of promising versus insufficient approaches, clarification of the terminology landscape, and definition of the initial research track.
- Established the scientific framing for the problem, including research hypotheses, methodological boundaries, known limitations, evaluation logic, and success criteria required for the work to be meaningful under real enterprise data constraints.
- Translated customer-specific language challenges into concrete technical research questions, data requirements, and benchmark design, creating a credible path toward more enterprise-adapted speech recognition without relying on large-scale domain audio collection.
- Defined which methods, papers, and experimental directions were worth re-implementing, adapting, or testing, and maintained technical oversight throughout execution to keep the research aligned with practical usability goals.
- Established and technically led an external applied research collaboration with UNISINOS, coordinating a multi-party research setup involving an AI Scientist, two Linguists, approximately three Professors, and additional internal contributors from the company side.
- Mentored and guided external researchers while maintaining scientific and technical control over the collaboration, ensuring continuity between the original problem framing, the experimentation path, and the final evaluation outcomes.
- Co-designed a novel multilingual open-vocabulary keyword spotting (OVKWS) architecture for the prompt-selection task, avoiding dependence on domain-specific keyword training data by using TTS-generated synthetic audio for arbitrary keywords.
- Co-designed an approach in which synthetic keyword audio is compared against the input audio stream through a cross-attention mechanism that learns acoustic similarity, improving flexibility for specialised vocabulary recognition and terminology-aware speech handling.
- Contributed to a four-stage incremental training strategy and hard-negative sampling approach to reduce false positives and improve robustness in specialised-vocabulary recognition scenarios.
- Compiled, adapted, and strengthened evaluation benchmarks and metrics for specialised-vocabulary recognition in enterprise speech settings, ensuring the research could be assessed against stakeholder-relevant success criteria.

Outcomes / Strategic Impact

- Created a credible and technically grounded path toward enterprise-adapted speech recognition for customer-specific terminology and specialised jargon under realistic low-data conditions.

- Addressed a core barrier to enterprise Speech AI adoption by advancing approaches to terminology injection and open-vocabulary handling that do not rely on large domain-specific fine-tuning datasets.
- Improved enterprise speech-to-text quality on company-domain benchmarks, with gains of up to **+25% Mean Reciprocal Rank (MRR)** in English and up to **+133% MRR** in Portuguese on selected benchmark evaluations.
- Expanded innovation capacity through a technically directed academic collaboration while keeping the work tightly anchored to enterprise deployment realities, evaluation rigor, and practical usability.
- Strengthened external research credibility through an accepted **LREC 2026** paper, **"A Dataset for Evaluating ASR on Specialised Vocabulary" (Submission 1226)**, and deepened internal capability in specialised vocabulary evaluation, terminology-aware recognition, and data-efficient Speech AI methods.
- Helped drive the work to publication-grade technical maturity, including a second full paper draft covering the system architecture, training strategy, and experimental design for the specialised-vocabulary recognition approach; the material remained internal due to company confidentiality constraints on external publication.

Ownership

Owned the research direction end-to-end from the company side: identified the enterprise blocker, established the scientific framing, qualified the literature and experimental path, set evaluation logic and success criteria, established the external collaboration, and maintained technical oversight throughout execution to steer the work toward usable, externally validated, and internally strategic outcomes.

Enterprise Speech AI Evaluation, Benchmarking and Stabilisation Frameworks

Period: 2024–2026 | Role: AI Scientist / Speech AI Evaluation Lead / Quality-Improvement Architect

Mission. Built the enterprise Speech AI evaluation and quality-improvement architecture from the ground up, enabling systematic benchmarking, diagnosis, and optimisation across multilingual, multi-domain, and customer-relevant speech scenarios.

Key Contributions

- Built the evaluation framework end-to-end as the primary AI owner, defining the methodology, benchmark structure, quality criteria, metrics, dashboards, and improvement logic for enterprise Speech AI.
- Assembled and curated a large-scale benchmarking asset spanning **~90 open-source datasets**, SAP data, and a planned expansion with commercial datasets, creating a reusable evaluation foundation across **~12 languages**, with scaling planned beyond **40**.
- Structured the benchmark across **100+ evaluation dimensions**, including speech modality, spontaneous / semi-spontaneous / read speech, gender balance, accents, thematic domains, audio length, noise conditions, products, languages, and interaction types such as human-to-human, human-to-machine, and machine-to-machine.
- Built automated evaluation pipelines for speech-to-text quality assessment across models, pipeline variants, input conditions, and product-relevant scenarios including online meetings, subtitling, translated transcription, marketing and learning content, and speech interfaces for agents and robotics.
- Integrated automated metrics with human evaluation and correlation analysis, making quality assessment more reliable, comparable, and actionable across models and data conditions.
- Led evaluation-driven improvement work for SAP-specific Speech AI across speech-to-text, Text-to-Speech, terminology adoption, phoneme-to-grapheme and spelling normalisation, language-specific switching, pause and noise handling, and vocoder-supported output consistency.

- Prioritised and validated multiple non-fine-tuning improvement directions, including data-specific pre-processing, terminology injection, prompt-based optimisation where model capabilities allowed, and post-processing strategies for offline transcription and subtitling scenarios.
- Designed patented stabilisation concepts for speech-to-text pipelines aimed at reducing hallucinations and skipped content through protected processing and verification logic.
- Designed decision logic that linked offline metrics, human evaluation, robustness checks, and improvement priorisation into a reusable enterprise quality-governance layer.
- Operated as the sole AI specialist on the initiative, setting direction from concept design and evaluation methodology through implementation guidance, benchmark creation, and improvement prioritisation, supported by software engineers.

Outcomes / Strategic Impact

- Established the core benchmarking and evaluation backbone for enterprise Speech AI, turning fragmented quality work into a reusable measurement and improvement system.
- Enabled continuous comparison of internal and commercial speech models across multilingual, multi-domain, and real-world scenario conditions.
- Delivered **up to 20% Word Error Rate improvement** in relevant enterprise speech-recognition scenarios through evaluation-driven adaptation, pre-processing, terminology handling, and post-processing.
- Improved robustness in noisy and latency-sensitive workflows and created a stronger basis for systematic optimisation of customer-facing speech use cases.
- Drove strong quality gains in Text-to-Speech through terminology-aware pronunciation control, language-specific normalisation, and output-quality improvements.
- Co-authored patented innovation through **"Adaptive Fidelity Pipeline for Minimizing Hallucinations and Skipped Content in Speech-to-Text Systems"** *(pending, 250089US01, Aug 18, 2025)*.

Ownership Owned the enterprise Speech AI evaluation and quality-improvement direction end-to-end as the primary AI lead: defined the methodology, built the benchmark and evaluation system, guided implementation, prioritised improvement paths, and translated Speech AI experimentation into measurable, reusable, enterprise-usable quality gains.

Japanese Enterprise TTS Optimisation: Terminology and Pronunciation Quality

Period: 2024–2025 | Role: Applied AI Research Lead / Mentor

Mission. Owned a Japan-focused applied AI direction addressing a priority enterprise TTS blocker: natural and accurate rendering of technical terminology, anglicisms, and abbreviations in Japanese in settings where extensive fine-tuning was impractical due to data sensitivity, compliance overhead, and cost.

Key Contributions

- Established Japanese TTS quality as a priority enterprise issue in terminology-heavy communication scenarios and characterised the main failure modes, including English-style rendering of Latin-script terms, abbreviation misinterpretation, inconsistent pronunciation, and degradation on longer term-dense sequences.
- Qualified feasible improvement paths through literature-based method review and targeted experimentation, selecting a lower-cost pipeline-based direction over broad fine-tuning, terminology-scale phonetic annotation, or architecture-heavy intervention.
- Mentored a student in Japan through sustained hands-on guidance on experimentation, corner-case analysis, evaluation data development, and validation with local language experts.

- Directed terminology- and pronunciation-focused quality improvement for Japanese TTS under real enterprise data and compliance constraints.

Outcomes / Strategic Impact

- Established a viable, lower-cost path to improve Japanese TTS quality without requiring extensive new audio-data creation or broad fine-tuning effort.
- Delivered early gains of **up to 15% Word Error Rate improvement** in relevant evaluation settings.
- Co-authored the patent **"Enhancing Pronunciation Accuracy and Naturalness in Japanese Text-to-Speech Systems without Extensive Fine-tuning"** (under review).

Ownership

Owned the problem framing, technical direction, and experimental path end-to-end; mentoring extended execution capacity, while the scientific direction and solution strategy remained defined and driven from your side.

4. AI Architecture, Data Foundations, and Future Systems

SAP Internal Speech Data Product Architecture

Period: 2024–2026 | Role: Speech Data Product Architect / Applied Researcher

Mission. Defined the architecture of a SAP-wide internal multilingual speech data product to turn fragmented, highly sensitive audio sources into governed, reusable enterprise AI assets for SAP-specific speech use cases.

Key Contributions

- Architected the internal speech data product end-to-end, defining how SAP-specific multilingual audio assets should be created, transformed, stored, catalogued, and reused across enterprise AI workflows.
- Designed the core data model across raw audio, segmented audio, transcripts, translations, speaker and session entities, terminology annotations, synthetic data, and processing artefacts, including the entity relationships, schema logic, constraints, and ownership structure required for traceable reuse.
- Defined the architecture to support synthetic-data generation and augmentation workflows, including TTS-based data creation with human-in-the-loop validation, as part of the broader multilingual speech data product design.
- Defined a **20+ stage** pipeline architecture spanning ingestion, segmentation, alignment, preprocessing, transformation, metadata engineering, validation, deduplication, licensing checks, and dataset assembly.
- Designed the storage architecture using **AWS S3** for raw and intermediate layers and **SAP DataSphere** as the governed target layer for downstream access, discovery, and consumption.
- Built the metadata and governance model across **100+ dimensions**, covering speech modality, interaction type, language, domain, sensitivity, consent, and derived processing metadata needed for compliant enterprise use.
- Structured the architecture for scale to **~40 languages** and downstream use across **STT, TTS, speech-to-speech translation, evaluation, subtitling, voice-over, meeting transcription, robotics, and agentic AI**.

Outcomes / Strategic Impact

- Created the architectural basis for a SAP-wide internal speech data product rather than isolated, one-off speech datasets.
- Established the schema, metadata system, and pipeline design needed to convert scattered and weakly governed audio sources into reusable enterprise speech data.
- Enabled future SAP-specific Speech AI adaptation and evaluation by defining a data product that supports multilingual reuse across research, benchmarking, and product-facing workflows.

- Positioned speech data as a governed internal asset class for enterprise AI rather than a fragmented project dependency.
- Established and led a cross-functional Speech Data Circle to align stakeholders on audio-data standards, requirements, process maturity, and reusable practices needed to operationalise the SAP internal speech data product.

Ownership: Owned the architecture end-to-end: defined the data product concept, designed the schema and entity model, set the storage and pipeline architecture, and translated speech-data complexity into an implementable enterprise design for compliant multilingual AI use.

Speech AI for Robotics and Voice-Based Intelligent Systems

Period: 2025–2026 | Role: Speech AI Research Contributor / Applied Innovation Contributor
Mission. Contributed to international applied research on Speech AI for robotics and voice-enabled intelligent systems, focusing on SAP-side speech capabilities for multilingual, deployment-constrained human–machine interaction scenarios.
Key Contributions

- Contributed to two applied research initiatives exploring speech input and output for robotics-related systems, including future warehouse and intelligent automation scenarios.
- Supported SAP-side work on speech model selection, domain adaptation, evaluation logic, and quality requirements for multilingual and noise-affected real-world conditions.
- Contributed to early research and optimisation work for smaller speech models, edge-device feasibility, quantisation-related optimisation, and deployment constraints including latency, memory, and offline capability.
- Worked within a large international collaboration spanning **SAP Japan**, **SAP Research & Innovation in Potsdam**, customer partner **TRUSCO Nakayama**, robotics partner **AIMBO Robotics**, and an industry-academia collaboration with **Nagoya University's Kawaguchi Laboratory**.
- Contributed speech-AI expertise on the SAP side for future scenarios in which speech interfaces support embodied AI, warehouse robotics, and intelligent enterprise systems connected to **S/4HANA** workflows.

Outcomes / Strategic Impact

- Extended applied research exposure into embodied AI, enterprise robotics, and future speech-enabled intelligent systems.
- Contributed to early feasibility work for SAP-side speech capabilities in robotics scenarios with real deployment constraints.
- Strengthened frontier-facing research signal through participation in a high-profile international collaboration spanning SAP, industry, customer, robotics, and academic partners.

Ownership
Contributed SAP-side speech-AI expertise in a collaborative, in-progress research setting, with focus on model selection, evaluation logic, domain adaptation, and deployment-aware speech quality considerations rather than end-to-end ownership of the broader robotics platform.

AI Architecture for Terminology-Aware Enterprise Use Cases

Period: 2023–2025 | Role: AI Architect / Strategic AI Advisor
Mission. Provided short-cycle architecture and feasibility input for selected enterprise AI use cases built around SAP's multilingual terminology assets.
Key Contributions

- Assessed architecture options for SAPTerm-based use cases involving multilingual terminology disambiguation and support-answer retrieval.
- Evaluated how LLM systems, RAG pipelines, knowledge graphs, and agentic workflow patterns could be applied credibly under enterprise constraints.
- Translated early-stage AI ideas into technically grounded solution directions and feasibility considerations for stakeholder discussions.

Outcomes / Strategic Impact

- Improved the technical quality of early-stage decision-making around selected SAPTerm-related AI opportunities.
- Helped assess future-facing terminology-aware AI use cases before deeper investment or product commitment.

Ownership

Provided architecture and feasibility judgment for selected terminology-centric AI ideas, turning early use-case concepts into more credible technical options for further discussion.

Agentic AI, Query Reduction, RAG, KGs, and Knowledge-Enhanced Enterprise AI

Period: 2024–2026 | Role: Applied AI Researcher / AI Architect / Innovation Advisor

Mission. Provided AI architecture, feasibility, and evaluation guidance for selected GenAI and hybrid-AI initiatives spanning query reduction, agentic workflows, RAG, and knowledge-enhanced enterprise use cases.

Key Contributions

- Led the AI-side feasibility and architecture direction for a **Query Reduction** initiative with SAP supplier **GlobalizeMe**, assessing how RAG, embeddings, and fine-tuning could reduce query volume and processing complexity in SAP language-service workflows involving ambiguous and high-volume enterprise data.
- Guided the supplier-led POC through architecture workshops, pre-research, technical direction-setting, and outcome review, establishing a basis for continued internal work under defined budget, fine-tuning, and data conditions.
- Acted as the principal AI advisor for a follow-on internal **Agentic AI POC** that extended the Query Reduction direction into more advanced workflow automation. Defined architecture, evaluation and control logic, reliability criteria, data requirements, tool-selection boundaries, and scenario design for multi-step enterprise workflows. Helped frame when retrieval, generation, tool use, and workflow orchestration were technically justified, and where agentic complexity would add risk without sufficient enterprise value.
- Provided conceptual advisory, architecture review, and feasibility judgment for **knowledge graph** and **RAG** directions built around SAP terminology assets.
- Helped screen and shape a portfolio of **100+ internal AI automation workflow ideas**, contributing AI feasibility input, architecture judgment, and prioritisation perspective; a condensed version of this portfolio was later reviewed in a leadership offsite in Brazil where you participated as the invited AI expert.

Outcomes / Strategic Impact.

- Improved technical decision quality around selected GenAI initiatives before deeper delivery commitment.
- Helped move Query Reduction from supplier-led feasibility work into a follow-on internal Agentic AI POC.

- Kept agentic AI, RAG, and knowledge-enhanced designs grounded in enterprise constraints, evaluation discipline, and realistic implementation paths rather than trend-driven experimentation alone.

Ownership:
Owned feasibility and technical judgment across selected initiatives, defining credible architecture paths, evaluation expectations, and continuation criteria for supplier-led and internal GenAI exploration.

5. Enterprise AI Advisory, Governance, and Organisational Enablement

AI Advisory to Leadership and Technical Stakeholders

Period: Ongoing, 2024–2026 | Role: Strategic AI Advisor / Technical Decision Partner
Mission. Provided principal-level AI guidance to leadership, product, and engineering stakeholders, helping shape technically credible decisions across fast-moving AI topics under real business, platform, and delivery constraints.
Key Contributions

- Advised leadership, product, and technical stakeholders on the feasibility, risks, and practical value of AI initiatives across **Speech AI, LLMs, RAG, knowledge graphs, agentic AI**, and related applied-AI directions.
- Helped decision-makers distinguish between trend-driven ideas and enterprise-credible AI opportunities by grounding discussions in model capability, data readiness, evaluation quality, architecture fit, and delivery realism.
- Contributed AI screening, architecture judgment, and prioritisation input across a portfolio of **100+ internal AI automation workflow ideas** aimed at reducing recurring manual effort and improving process efficiency.
- Participated as the invited AI expert in a leadership offsite in **Brazil**, where a condensed portfolio of future AI automation opportunities was reviewed from an enterprise AI feasibility perspective.
- Prevented weak AI directions by challenging unsuitable model choices, inadequate data or evaluation setups, misaligned agentic designs, and poor storage or data-architecture assumptions.

Outcomes / Strategic Impact

- Improved decision quality around AI investments, technical direction, and experimentation priorities.
- Reduced the risk of low-feasibility or poorly framed AI initiatives before deeper delivery commitment.
- Strengthened internal decision-making by providing a high-trust translation layer between emerging AI patterns and enterprise execution constraints.
- Helped keep advanced AI exploration tied to realistic automation, cost-efficiency, and implementation-value scenarios rather than trend-driven experimentation alone.

Ownership
Acted as a principal-level technical advisor, converting fast-moving AI developments into grounded architectural and strategic decisions across ambiguous, high-impact enterprise contexts.

Enterprise AI Governance, Stakeholder Alignment, and Execution Readiness

Period: 2024-2026 | Role: Cross-Functional AI Execution Lead / Enterprise AI Advisor
Mission. Enabled enterprise AI initiatives to move forward under legal, licensing, data-protection, procurement, and supplier constraints by converting governance complexity into executable technical paths.
Key Contributions

- Led the cross-functional execution layer required to move AI initiatives through governance, licensing, procurement, supplier, and stakeholder constraints in enterprise environments.
- Translated legal, operational, and business constraints into workable technical paths, reducing the gap between AI ambition and implementation feasibility.
- De-risked AI initiatives by surfacing organisational blockers early and aligning stakeholders around viable execution approaches.
- Guided procurement and vendor-readiness processes for AI-relevant data assets, aligning budget, legal review, supplier onboarding, licensing, and use-case-specific governance requirements.
- Defined concrete handling processes for highly sensitive AI data, including compliant use of open-source, commercial, and enterprise-internal datasets in production AI contexts.
- Established a cross-functional alignment mechanism for speech-data standards and operational readiness, improving organisational preparedness for SAP-specific Speech AI initiatives.

- Established and led the **Speech Data Circle** as a cross-functional mechanism for aligning stakeholders on speech-data standards, requirements, and operational readiness, reducing fragmentation in audio-data practices and supporting SAP-specific Speech AI execution.

Outcomes / Strategic Impact:

- Reduced the risk of technically strong but organisationally non-viable AI initiatives.
- Improved execution readiness across technical teams, suppliers, central functions, and governance stakeholders.
- Strengthened the organisation's ability to operationalise AI under real legal, licensing, procurement, and data-protection constraints.
- Defined practical processes and playbooks for scaling highly sensitive enterprise audio data in production AI environments.
- Improved organisational preparedness for SAP-specific Speech AI through the **Speech Data Circle**, creating a reusable cross-functional alignment mechanism around audio-data standards and execution readiness.

Ownership Owned the cross-functional execution layer required to move AI initiatives through governance, licensing, procurement, and stakeholder constraints.

AI Capability Building, Knowledge Distribution and Organisational AI Maturity

Period: Ongoing, especially 2024–2026 | Role: AI Capability Lead / Organisational Enablement Lead

Mission. Strengthened organisational AI maturity by designing and delivering structured enablement that translated advanced AI topics into practical understanding, aligned expectations, and greater readiness across technical, business, and domain-specific audiences.

Key Contributions

- Designed and co-delivered structured AI enablement programs, including multi-day sessions for **LXE/LXT stakeholders in Germany and Brazil** and focused workshops for translator-oriented groups.
- Adapted advanced AI content to different audience needs, making topics such as **agentic AI, knowledge graphs, GenAI direction, and applied AI strategy** more actionable for specialist and non-specialist stakeholders.
- Helped build shared language and stronger alignment across teams by turning fragmented AI understanding into more practical, usable organisational knowledge.
- Supported colleagues in understanding how to adapt, re-skill, and contribute to broader AI transformation as AI capabilities expanded across the organisation.

- Acted as an informal AI strategy ambassador across teams.

Outcomes / Strategic Impact

- Increased practical AI maturity across mixed stakeholder groups.
- Improved cross-functional understanding of AI opportunities, risks, and implementation constraints.
- Strengthened the organisational foundation for more aligned, informed, and realistic AI adoption across teams and geographies.

Ownership

Owned a meaningful part of the organisation's AI capability-building layer by translating fast-moving AI topics into structured enablement, shared understanding, and greater operational readiness.

6. Core Applied AI, Conversational, Retrieval, and NLP Foundation

Enterprise Multilingual Retrieval, Translation AI, and Evaluation Architecture

Period: 2019–2024 | Role: Applied AI Scientist / NLP Architect / Retrieval & Evaluation Systems Lead

Mission. Served as the primary NLP / ML / AI specialist in a lean cross-functional team rebuilding the next-generation **SAP Translation Hub (STH) Multilingual Text Repository (MLTR)**, a large-scale enterprise multilingual retrieval platform for thousands of SAP customers operating on **millions of verified translations** across **40+ target languages** and **2,000+ language combinations**.
Owned the AI-facing design of retrieval logic, design and implementation of the evaluation frameworks, quality scoring, and improvement methodology for enterprise translation and localisation workflows.

Key Contributions

- Co-defined and helped drive the transition from legacy keyword-based retrieval to semantic retrieval in a large-scale multilingual enterprise repository, shaping retrieval logic, internal retrieval formats, and AI-relevant comparison methodology.
- Designed and implemented the automated evaluation architecture for translation retrieval, neural machine translation, and proposal-quality comparison, creating a repeatable measurement layer for system evolution, benchmarking, and decision support.
- Implemented a multidimensional evaluation framework combining retrieval, similarity, and translation-quality metrics - including edit-distance-based indicators, retrieval coverage, similarity scores, quality correlations, BLEU, COMET, ChaacTER, and embedding-based similarity - within scalable benchmarking pipelines generating 36+ performance reports per cycle and enabling continuous quality monitoring.
- Engineered automated benchmarking workflows across language pairs, query types, search configurations, and proposal scenarios, significantly increasing the speed, consistency, and depth of comparative analysis.
- Created representative test sets and behaviour-oriented evaluation scenarios simulating real translator usage patterns, aligning offline evaluation more closely with downstream productivity and business value.
- Led applied research and experimentation to improve retrieval relevance and proposal usefulness for both short-text and long-text enterprise translation scenarios, including large-scale application contexts involving 25,000+ custom objects.
- Designed and built quality and confidence scoring approaches for retrieval outputs, supporting safer automation, stronger ranking decisions, and improved proposal selection in downstream workflows.

- Developed Python-based evaluation services, data exports, and ML/NLP prototypes supporting retrieval analysis, proposal assessment, and translation-quality benchmarking.

Outcomes / Strategic Impact

- Improved proposal quality by up to 10% in cross-domain retrieval scenarios.
- Established a scalable evaluation and benchmarking foundation that made retrieval and translation quality more measurable, comparable, and easier to improve over time.
- Strengthened the AI and evaluation backbone for multilingual enterprise language services used in large-scale customer rollouts.
- Enabled more evidence-based retrieval and proposal-quality decisions by replacing ad hoc comparisons with repeatable benchmarking and quality measurement.
- Contributed to the technical foundation behind a platform used to translate thousands of strings at a time, adopted in large-scale projects by Fortune Global 500 organisations deploying SAP worldwide. Its customer case studies, including Siemens Healthineers and Daimler Truck, reported up to 93% reduction in manual translation effort, 99%+ proposal accuracy, 87% faster translation timelines, and 1,250 hours of developer effort saved.

Ownership:
Acted as the primary owner of the NLP / ML / AI dimension of the system: retrieval relevance, evaluation design, quality measurement, scoring logic, experimentation, and improvement direction.

Cross-Component Data Architecture and Interoperability

Period: 2019–2024 | Role: Cross-Team Technical Lead / Data Architecture Driver
Mission. Improved cross-component data architecture across core LX systems by defining more robust interchange formats, shared processing logic, and clearer communication patterns for enterprise language workflows.

Key Contributions

- Led the cross-team **Data Formats** group focused on **JSON**, **XLIFF**, and **Box-in-the-Box** approaches for structured data exchange across core LX components.
- Defined and refined internal interchange formats used across **five core components**, improving consistency and clarity in how linguistic and processing data moved between services.
- Designed and implemented a shared library for processing the common interchange format, reducing duplicated handling logic across components.
- Oversaw retrieval-related data structures and communication mechanisms, improving interoperability across data-processing pipelines.
- Drove standardisation of cross-service data movement, helping make enterprise ML/NLP workflows more maintainable and extensible.

Outcomes / Strategic Impact

- Reduced friction in communication between systems and teams by creating a more standardised cross-component data-exchange model.
- Improved maintainability and scalability of AI-related workflows through shared processing logic and clearer data contracts.
- Lowered communication overhead and created a more extensible foundation for large-scale linguistic and enterprise data exchange.
- Demonstrated platform-level ownership beyond isolated model development.

Ownership Led cross-team alignment on data formats and interoperability, defining common exchange patterns and shared processing mechanisms across core technical components.

Terminology Intelligence, Linguistic Data Quality, and Patentable Language Innovation

Period: 2019–2025 | Role: Applied Researcher / Data Scientist / Inventor

Mission. Advanced enterprise language intelligence by combining large-scale linguistic data analysis, terminology-focused NLP research, metadata quality work, and patentable innovation for more precise and reusable language processing.

Key Contributions

- Conducted extensive analysis of large-scale linguistic and translation-related datasets.
- Conducted large-scale analysis of SAP linguistic and translation data, including preparation and assessment of an inventory covering **12+ million multilingual term pairs** and a broad metadata landscape.
- Identified data issues that limited AI usability, including design inconsistencies, missing bilingualism, outdated synonyms and releases, unclear ambiguities, abbreviation-related problems, and metadata inconsistencies across terminology assets.
- Improved linguistic data quality and reliability through metadata diagnosis, quality correction, and categorisation approaches that made complex data collections more structured and usable.
- Built data-driven foundations for stronger terminology handling, ambiguity reduction, and semantic-domain consolidation across SAP translation data.
- Contributed to patentable innovation in **semantic domain assignment**, **terminology disambiguation**, and **abbreviation detection with mapping to original full terms**, connecting practical enterprise language problems with formal IP creation.

Outcomes / Strategic Impact

- Improved the reliability and reusability of enterprise language data for downstream AI use cases by strengthening metadata quality, data structure, and analytical visibility into failure points.
- Prepared a data-driven basis for consolidating and reducing semantic domains across SAP translation data, improving harmonisation and reuse for enterprise AI scenarios.
- Strengthened precision in enterprise language processing through reusable innovation assets with strategic IP value.
- Co-authored the issued patents including **"Semantic Domain Assignment Referencing Governance Domains and Term Databases"** *(US-12518105-B2, 6 January 2026)*, **"System and Method Performing Terminology Disambiguation"** *(US-12386820-B2, 13 August 2025)*, and **"Detection of Abbreviation and Mapping to Full Original Term"** *(US-12067370-B2, 20 August 2024)*.

Ownership

Connected real enterprise language-data problems with analytical diagnosis, terminology-focused NLP research, and patentable innovation, turning large-scale linguistic data work into reusable AI foundations and strategic IP.

7. Earlier Industry and Pre-AI Leadership Foundation

Enterprise Conversational AI for Technical Support at SAP

Period: 2018–2019 | Role: End-to-End Applied AI Developer / Data Scientist / Computational Linguist

Mission. Designed and developed an ML-based conversational agent for technical support, combining dialogue management, semantic search, and practical system deployment in an enterprise context.

Key Contributions

- Built a task-oriented conversational AI system from scratch as part of thesis-driven work on dialogue management and deep learning for support-oriented agents.

- Combined dialogue management with semantic search to improve relevance and effectiveness in technical-support interactions.
- Owned the full delivery path across data preparation, filtering, feature setup, model exploration, system design, deployment, and iterative improvement.
- Coordinated working students supporting testing and web-design-related implementation tasks.

Outcomes / Strategic Impact

- Prototyped a deep-learning chatbot for internal enterprise technical support use.
- Demonstrated early end-to-end ownership across research, experimentation, implementation, and deployment.
- Established a strong practical foundation for later work in applied NLP, enterprise AI systems, and Speech AI.

Ownership Owned the system end-to-end as an early applied AI builder, spanning research, development, deployment, and practical coordination.

NLP / AI Projects at Empolis Information Management GmbH

Period: 2018 | Role: Data Scientist / Software Developer in Machine Learning and Computational Linguistics

Mission. Contributed to applied NLP and text-analytics projects in a smart-data environment, with focus on information extraction, scientific search, and domain-specific language analysis.

Key Contributions

- Worked on information extraction, text mining, and linguistic-analysis models in practical industry settings.
- Supported development of an NLP-driven scientific paper search engine in the pharmaceutical domain.
- Contributed to sentiment-analysis work as well as cross-functional testing and integration activities.

Outcomes / Strategic Impact

- Strengthened early practical expertise in applied NLP, text mining, and domain-specific information retrieval.
- Expanded hands-on experience in translating language-technology methods into usable analytics and search-oriented solutions.

Ownership

Contributed to industry NLP and ML delivery across modelling, implementation, and integration activities.

Spiegel Institut

Period: 2016 | Role: Computational Linguist

Mission. Contributed to an early spoken-language project in the context of autonomous driving and human–machine interaction for vehicle communication systems.

Key Contributions

- Worked on NLP processing of spoken data for a vehicle communication framework.
- Applied computational-linguistics methods in an automotive-related HMI setting.

Outcomes / Strategic Impact

- Built early domain experience in spoken language processing and AI-related mobility scenarios.
- Broadened applied NLP exposure into human–machine interaction contexts.

Ownership
Contributed language-technology expertise within a future-facing automotive communication project.

Leadership and Entrepreneurial Experience Before Full AI Focus

Period: 2006–2015 | Role: Property Expert / Founding Partner and CEO / Head of Quintessentially Estates Ukraine
Mission. Built business, leadership, analytical, and decision-making experience in international and entrepreneurial environments before transitioning fully into AI-focused roles.
Key Contributions

- Worked in analytical and client-facing business environments with strong commercial and stakeholder orientation.
- Led a company as CEO and held leadership responsibility in an international real-estate context.
- Built durable foundations in strategic thinking, ownership, business positioning, and decision-making under real-world constraints.

Outcomes / Strategic Impact

- Added a strong business and leadership dimension to AI profile.
- Strengthened later credibility in advisory, stakeholder-facing, and cross-functional leadership roles beyond purely technical execution.

Ownership
Held real leadership and entrepreneurial responsibility before entering full-time AI work, building the business judgment and executive maturity that now strengthen senior technical and advisory roles.