



www.spitch.ch

AA Agent Assist Whitepaper

Every Voice Matters. Every Insight Counts.



TABLE OF CONTENTS

1	Spitch Agent Assist in day-to-day customer service	3
2	How Spitch Agent Assist supports agents during the conversation	4
3	What the University of St. Gallen investigated	5
4	Piloting: prove the impact before scaling	6
5	From pilot to permanent operation	7
	Source	8

01 MINI-PAPER

1. SPITCH AGENT ASSIST IN DAY-TO-DAY CUSTOMER SERVICE

What real-time support delivers – and how findings from the University of St. Gallen are translated into operations

Key message. Spitch Agent Assist supports customer service employees during customer conversations with relevant knowledge, concrete guidance, and configurable quality prompts. The solution is not meant to replace people, but to make them more capable of acting in demanding situations.

Why real-time support is gaining importance

Automation is increasingly handling simple, standardized requests. As a result, the conversations that remain for employees are often more complex, more emotional, or higher-risk. At the same time, agents have to listen attentively, look up information in several systems, give correct expert advice, comply with process requirements, and document the conversation. Doing all of this at once creates cognitive load, especially when the customer is impatient, annoyed, or uncertain.

What is Spitch Agent Assist?

Spitch Agent Assist is AI-based real-time support for customer service employees. The solution analyzes the ongoing conversation context and provides suitable support directly within the agent desktop. Depending on the use case and configuration, it can:

- **Provide knowledge:** display relevant articles, product information or process content from approved sources.
- **Give orientation:** surface next steps, conversation guides or situation-specific recommended actions.
- **Safeguard quality:** act as a reminder of mandatory information, process steps or defined quality criteria.
- **Prepare after-call work:** generate conversation summaries and structured documentation content.

Technically, Spitch Agent Assist is one of several modules of the Spitch Collaborative Agentic AI Platform. How this modular architecture can be used operationally – for example, by connecting it with Knowledge Management, Speech Analytics, Quality Management, or Agent Training – is placed in the context of the learning loop in Chapter 5.

Why the University of St. Gallen research matters

A joint investigation with the University of St. Gallen examined whether – and under which conditions – this form of support genuinely eases the workload. The study thus complements the technical product perspective with independent scientific evidence. This paper first explains the solution, then summarizes the University of St. Gallen findings in accessible terms, and derives an operational approach from them.

02 SOLUTION

2. HOW SPITCH AGENT ASSIST SUPPORTS AGENTS DURING THE CONVERSATION

The key difference from a conventional knowledge search lies in context awareness: employees do not have to trigger the support entirely themselves. Agent Assist processes the ongoing conversation, detects relevant signals, and brings suitable content into the working context. The specific scope of functionality is configured per call reason, role and process.

Moment	Possible support	Operational objective
Before the conversation	Customer and case context, previous contacts, relevant permissions	Shorten preparation and understand the starting situation
During the conversation	Knowledge recommendations, guidance, next steps, mandatory and quality prompts	Reduce searching, increase confidence and maintain conversation flow
After the conversation	Summary, classification, and prepared documentation	Speed up after-call work and improve completeness
Across conversations	Usage, quality and outcome data for analytics, QM, and training	Continuously improve knowledge, processes, and skills

People remain in control

Agent Assist makes recommendations, but does not automatically make expert decisions on the employee's behalf. The agent assesses the prompt in the context of the conversation and can accept, ignore, or correct it. For this to work, the source, currency, and validity of the underlying content must be transparent.

Embedded support instead of another tool

The solution can be connected to telephony and contact center platforms, CRM systems, and company-specific knowledge sources. The goal is support within the existing workflow. Additional context switching would increase cognitive load and therefore work against the actual purpose.

Design principle. The greatest benefit does not come from generating as many prompts as possible, but from a small number of well-founded recommendations at the right moment.

03 SCIENTIFIC EVIDENCE

3. WHAT THE UNIVERSITY OF ST. GALLEN INVESTIGATED

The University of St. Gallen white paper examines how AI-based assistance can support customer service employees under pressure. The research combines three perspectives: a field trial with Spitch Agent Assist in real service settings at a bank and a health insurer, an online study on task completion, and a laboratory study on cognitive load. The investigation was carried out independently by researchers at the University of St. Gallen.[1]

The key findings

Finding of the University of St. Gallen research	Implication for operational deployment
Relevant information can be available faster; searching and hold time can be reduced.[1]	The best fit is knowledge-intensive requests with long search time or many system switches.
Support can reduce cognitive load, particularly in difficult or impolite conversations.[1]	Prompts must be short, prioritized, and as undistracting as possible.
New employees have a particularly high potential benefit.[1]	Agent Assist can support onboarding, but it does not replace expert qualification.
The knowledge base, interface design, and human control shape effectiveness.[1]	Knowledge ownership, approval processes, and feedback channels are part of the solution, not downstream details.

How to interpret the results

The research provides a robust rationale for real-time support and identifies the relevant mechanisms at work. It does not, however, promise a universal AHT reduction or blanket ROI. The economic effect depends, among other things, on the call reason, the quality of company knowledge, integration, and actual usage. The results therefore have to be operationalized and measured in a customer-specific pilot.

Bridge to operations. The study answers why Agent Assist can help. The pilot has to show where, how strongly, and under which conditions the solution takes effect in the particular organization.

04 MEASUREMENT MODEL

4. PILOTING: PROVE THE IMPACT BEFORE SCALING

A viable pilot does not start with the technology, but with a testable hypothesis. For example: “For complex contract inquiries, Agent Assist reduces search and hold time without degrading expert quality or customer satisfaction.” For onboarding, the hypothesis might be: “With Agent Assist, new employees reach a defined quality level faster.”

The pilot in five steps

1. **Select the use case.** A clearly delimited call reason with high knowledge or process complexity; not a company-wide rollout as the first step.
2. **Establish the baseline.** Measure search time, hold time, after-call work, expert quality, repeat contacts, and perceived workload before activation.
3. **Enable comparison.** As far as operationally feasible, compare results with and without assistance, or across comparable groups and time periods.
4. **Explain usage.** Do not only measure end results, but also analyze how individual recommendations are displayed, opened, accepted, ignored, and corrected.
5. **Decide go/no-go.** Scale only where there is demonstrable improvement, sufficient acceptance, and manageable quality and governance risks.

A balanced set of metrics

Dimension	Example metrics
Efficiency	search time · hold time · AHT · after-call work time · system switches
Quality	expert accuracy · FCR · process/compliance adherence · documentation quality
Employees	cognitive load · trust · adoption rate · time to confident handling
Customers	CSAT · customer effort · repeat contact · call abandonment
Assistance quality	relevance · timing · acceptance/dismissal rate · correction rate · missing content


Decision rule. No single KPI decides. A shorter AHT is not a success if errors, repeat contacts, or workload increase. What matters is the overall picture of efficiency, quality, acceptance, and customer outcome.

05 OPERATING MODEL & GOVERNANCE

5. FROM PILOT TO PERMANENT OPERATION

The closed improvement loop

Spitch Agent Assist should not be run as an isolated tool. Combined with Speech Analytics, Quality Management and Agent Training, it creates a continuous learning loop: the solution supports live interactions, the organization measures usage and outcomes, improves knowledge or process logic, and delivers targeted training for the skill gaps identified.



Responsibilities in day-to-day operations

- **Operations:** prioritizes use cases, defines target processes, and owns usage in day-to-day business.
- **Knowledge Owner:** manages sources, currency, approvals, and the correction of faulty content.
- **Quality & Training:** connects usage data with quality deviations and suitable learning measures.
- **IT, Data Protection, and Compliance:** own integration, permissions, logging and permissible use of data.
- **Product Owner Agent Assist:** brings together the backlog, KPIs, and cross-departmental decisions.

Prerequisites for scaling

Extending to further teams or call reasons makes sense once the pilot shows a clear net effect, knowledge and approval processes work reliably, the integration does not create new context switches, and employees understand both the purpose and the limits of the support. A cross-functional body should prioritize new use cases and regularly assess impact, risks, and acceptance.

Conclusion. The University of St. Gallen research shows why real-time support matters. Spitch Agent Assist translates this approach into everyday work: relevant knowledge, guidance, and quality prompts during the conversation – combined with a measurable learning and improvement process. Operational impact arises when technology, knowledge, processes, and responsibility are designed together.

Next step. The University of St. Gallen study shows where Agent Assist can be particularly relevant: in complex conversations, difficult situations, and when onboarding new employees. For organizations evaluating Agent Assist, a focused pilot with a clearly defined use case is a good starting point. It makes the operational impact measurable against concrete success criteria before a decision on broader rollout is taken.

SOURCE

[1] University of St. Gallen, Institute of Information Systems and Digital Business (2026): White Paper on Agent Assist. Independent investigation comprising field, online and laboratory studies; in the field trial, Spitch Agent Assist was observed in real service settings.