Victor Barres, PhD

AI Researcher · Conversational Agents & Evaluation

victor.barres@gmail.com · victorbarres.github.io · linkedin.com/in/victor-barres · Google Scholar

Experience

Sierra Sep 2024 – Present
AI Researcher San Francisco, CA

- Lead the **τ-Bench family** of agent benchmarks — research direction, codebase, public leaderboard (taubench.com), and ecosystem of extensions: τ^2-Bench (dual control), τ-Knowledge (knowledge retrieval), τ-Voice (full-duplex voice), and τ^3-Bench (the combined release with community-contributed task and code improvements). Three papers from the family accepted to **ICML 2026**.
- Kickstarted Sierra's internal **agent evaluation system**: set up the proof of concept and did the original work on user-simulation-driven stress testing for pre-release agent hardening across deployed customer-service domains.
- Organizer & judge of the Sierra τ^2-Bench Custom Track at the AgentX–AgentBeats competition (Berkeley RDI, Fall 2025 – Spring 2026).

Elemental Cognition Jun 2023 – Jul 2024
Senior Researcher (NLP) Remote

- Developed core algorithms for EC's hybrid AI architecture, integrating LLMs with formal reasoning engines.
- Led the integration of multi-round RAG with formal models of causation for knowledge-intensive multi-hop QA, published as *From Generating Answers to Building Explanations* at NAACL 2025 (Industry).
- Built conversational agents for two reasoning-heavy settings: a logic-programming assistant that elicits and translates user requirements into a logic program, and an ASP-reasoner-backed constraint-solving agent.

Uniphore.ai Dec 2019 – Jun 2023
Principal NLP Scientist & NLP Research Lead, previously Senior NLP Scientist (2019–2021) Palo Alto, CA

- Built and led Uniphore's NLP Lab (7+ scientists). Set short-, mid-, and long-term research strategy for understanding human–human spoken conversation in the contact-center domain.
- Shipped foundational and applied work across LLM benchmarking, prompt engineering, structured knowledge grounding, summarization, semantic parsing, dialogue-act and dialogue-flow modeling, and real-time caller-intent recognition.
- Kickstarted the team, deployed first NER / redaction / call-driver models, and stood up cloud infrastructure and CI/CD for production NLP.

Astound.ai Sep 2017 – Dec 2019
AI Scientist (NLP) Menlo Park, CA

- Core member of the team that built, from the ground up, Astound's domain-expert conversational agent for IT / HR / Finance enterprise self-service.
- Led the neuro-symbolic side of the agent's NLP stack — word/sentence embeddings, semantic parsing, knowledge-graph embeddings, and semantic inference — integrated with a modular dialogue manager (domain classification, intent recognition, search, small talk, sub-flows).
- Built NLP-side APIs over Astound's ontology and extended it to support inference-driven multi-turn clarification, establishing a separation of concerns between knowledge, lexicon, and grammar engineers.

Education

University of Southern California Sep 2010 – Aug 2017
PhD, Computational Cognitive Neuroscience Los Angeles, CA

Dissertation: *Schema Architecture for Language–Vision Interactions: A Computational Cognitive Neuroscience Model of Language Use.* Advisor: Michael A. Arbib.

École Normale Supérieure Jun 2008 – Jun 2010
MS, Cognitive Science Paris, France

École Polytechnique Jun 2003 – Jun 2006
MS, Physics Palaiseau, France

Lycée Louis-le-Grand Jun 2001 – Jun 2003
Classes Préparatoires – Mathematics & Physics Paris, France

Selected Publications

τ-Voice: Benchmarking Full-Duplex Voice Agents on Real-World Domains

V. Barres*, S. Ray*, V. Dhandhania*, K. Narasimhan. arXiv:2603.13686, 2026. Accepted to ICML 2026.

τ-Knowledge: Evaluating Conversational Agents over Unstructured Knowledge

Q. Shi, A. Zytek, P. Razavi, K. Narasimhan, V. Barres. arXiv:2603.04370, 2026. Accepted to ICML 2026.

τ^2-Bench: Evaluating Conversational Agents in a Dual-Control Environment

V. Barres*, H. Dong*, S. Ray, X. Si, K. Narasimhan. arXiv:2506.07982, 2025. Accepted to ICML 2026.

From Generating Answers to Building Explanations: Integrating Multi-Round RAG and Causal Modeling for Scientific QA

V. Barres, C. McFate, A. Kalyanpur, K. K. Saravanakumar, L. Moon, N. Seifu, A. Bautista-Castillo. NAACL 2025 (Industry)

LLM-ARC: Enhancing LLMs with an Automated Reasoning Critic

A. Kalyanpur, K. Saravanakumar, V. Barres, J. Chu-Carroll, D. Melville, D. Ferrucci. arXiv:2406.17663, 2024

Template Construction Grammar: From Visual Scene Description to Language Comprehension and Agrammatism

V. Barres, J. Lee. Neuroinformatics 12(1), 2014

Synthetic Event-Related Potentials: A Computational Bridge between Neurolinguistic Models and Experiments

V. Barres, A. Simons, M. Arbib. Neural Networks 37, 2013

* Equal contribution. Full list: victorbarres.github.io/publications/

Languages

English (fluent) · **French** (native) · **Spanish** (conversational)