Sebastián Bobadilla Suárez

PhD in Cognitive Neuroscience · VP, Data Science · Agentic AI, Knowledge Graphs & Computational Modelling

UK (Remote) | sebastian.bobadilla.s@gmail.com | bobaseb.github.io | github.com/bobaseb | in/sebastian-bobadilla-suarez

RESEARCH PROFILE

Computational cognitive neuroscientist and applied AI research leader bridging formal theory, statistical modelling, and production agentic AI systems. Author of 14 peer-reviewed articles and preprints spanning cognitive psychology, neuroimaging, behavioural economics, and machine learning (*Nature Communications*, *Nature*, *Communications Biology*, *Cognitive Psychology*, *Journal of Neuroscience*, *Journal of Risk and Uncertainty*, *arXiv*, *PsyArXiv*). Currently VP, Data Science in financial-services AI, leading ontology/knowledge-graph engineering and empirical research on agentic coding following postdoctoral and doctoral research at University College London and The Alan Turing Institute.

EDUCATION

PhD in Cognitive Neuroscience · University College London (UCL), London, UK — 2013 – 2018

Thesis: *The Interaction between Task Goals and the Representation of Choice Options in Decision-making*. Supervisors: Profs. Brad Love & Tali Sharot. Funded by CONACYT Doctoral Scholarship. Produced 4 peer-reviewed publications on decision heuristics, belief updating, delegation, and computational neuroimaging.

Enrichment Year Student (Data Science) · The Alan Turing Institute, London, UK — 2016 – 2017

UK National Institute for Data Science. Research on neural similarity measures and machine classifiers; masterclasses in machine learning and statistics.

MSc in Social Cognition: Research and Applications (Distinction) · University College London (UCL), London, UK — 2012 – 2013

Thesis: *Uncertainty, Confidence, and Suboptimality in Choice* (eye-tracking study of value-based decision confidence). Funded by CONACYT Scholarship.

BA in Social Anthropology (Honours, GPA 9.7/10) · Escuela Nacional de Antropología e Historia (ENAH), Mexico City — 2007 – 2011

Thesis: *Moral Hierarchies of Madness in the Popolucas of Soteapan, Veracruz* (ethnography of psychopathology in an indigenous community; 120+ days of fieldwork).

EXPERIENCE

OnCorps — AI agents for fund-operations oversight in financial services — UK (Remote)

VP, Data Science (from Feb 2026) · VP, Machine Learning (Dec 2025 – Feb 2026) — Dec 2025 – present

- **Ontology & Knowledge Graph Programme:** Started and drive the company's financial-services ontology programme (OWL/RDF/FIBO, Apache Jena Fuseki/SPARQL), authoring 95% of its human-written commits; exposed to product agents via read-only MCP with per-client data segregation and rule provenance. A dedicated ontology squad is now being formed, with this role as technical leader.
- **Production Anomaly-Detection Agents:** For a global fixed-income asset manager (∼$2.3 trillion AUM), evolved daily fund income-projection anomaly detection from MAD-score statistics to an agentic investigation layer. Onboarded 5 federated Snowflake feeds via Unity Catalog Lakehouse Federation, built 15 rule-based and 23 statistical checks, and encoded analyst playbooks as agent skills. Audited and fixed production defects (sign error in 107 of 121 exports, replay non-determinism, stale comparisons), freed 1.61 TB of leaked storage, and added fail-closed guards.
- **Contract-First Agent Platform:** Tech lead for self-service check authoring across 2 products. Designed a unified OpenAPI 3.1 agent platform with semver breaking-change gates, generated Python/TypeScript SDKs, offline evaluation suites (golden sets, deterministic scorers, LLM judges), and fail-closed security scans blocking destructive SQL and external URLs.
- **Always-On Multi-Agent POC (Legion):** Built an internal proof-of-concept always-on agent and alternative to Claude Code (not in production; currently in personal work use and championed for wider colleague adoption) for multi-week memory and multi-repo orchestration. Designed a hierarchical memory tree compressing multi-week conversations into 128 KB without deleting history, an orchestrator–manager–worker crew driving Claude Code across repositories, and web/TUI/iOS/chat interfaces; hardened via 4 security review rounds and mutation testing (1,095 passing tests).
- **Agentic Coding Transformation:** Led a company-wide agentic-coding programme for 2 months (Apr–Jun 2026; 20 epics across 2 tracks), then handed over. Trained all engineers, pushed security hooks and Ruff security lint gates, set a $200/day AI spend cap, launched a shared agent-skills plugin marketplace, and reported weekly to the ELT and board.
- **Delivery Telemetry & Empirical Research:** Analysed 6 months of engineering telemetry across 442 repositories, identifying distinct per-product bottlenecks (e.g., 27.9-day release waits vs. 40.6-day work-item lifetimes) to deliver

4 targeted CEO decisions. First author of *Audit the Scaffold, Not the Checkpoint* (arXiv:2609.34924), combining Lean 4 proofs, SWE-bench experiments, and multilevel modelling of 378 production agentic sessions.

- **Leadership, Governance & Agent Orchestration:** Manage 2 direct reports; interviewed and supported 2 data-science hires; reviewed 385 PRs in 2026; member of the ISO 42001 AIMS management board. Orchestrated parallel coding agents (600+ merged PRs across 19 repos in 2026) through human review gates and Red Zone security sign-off.

Sr. Director, Machine Learning (from Apr 2024) · Director, Machine Learning Apr 2022 – Dec 2025

- **Explainable Anomaly Detection:** Designed and shipped massively univariate MAD-score anomaly detection on Databricks for a ∼$2.3 trillion AUM fixed-income asset manager, ranking daily fund income-projection line items against 3-month histories and multiple temporal deltas in parallel.
- **Production Dashboards & Stack Modernisation:** Owned 4 production pricing dashboards for two Fortune 500 IT services firms; migrated 10,000+ LOC from R Shiny to Python Dash on a custom component library, modernised ETL and serving from Pandas/SQLite to Polars, Python 3.12, Dash 3, and AWS RDS with query caching, and deployed via Kubernetes.
- **Profit Modelling & Applied AI Systems:** Built a deal profitability model for the same two Fortune 500 IT services firms ($R^2 = 0.935$ and > 0.8), outperforming prior modelling efforts and re-engaging a key client. Built RAG and agent chatbots (AWS Bedrock, OpenAI, Qdrant/Databricks Vector Search, Langfuse/MLflow tracing, MCP) for a ∼$21B-revenue IT services firm, a SQL-agent prototype for one of the world's largest asset managers (∼$13 trillion AUM), and a 94%-accuracy SBERT classifier.
- **AI Tooling Standardisation:** Piloted and ran the risk assessment that established Claude Code and Cursor as standard AI coding tools; trained engineers and non-technical executives.

University College London (UCL) London, UK

Research Associate in Computational Modelling of Brain Imaging Data Jan 2018 – Nov 2021

- Conducted computational cognitive neuroscience research with Prof. Brad Love, applying statistical learning, regularized regression, neural similarity metrics, and deep learning to human fMRI, intracranial electrophysiological, and behavioural datasets.
- Published 4 peer-reviewed articles as first or senior-collaborating author (*Nature Communications*, *Cognitive Psychology*, *Communications Biology*, *Computational Brain & Behavior*), co-authored a *Nature* many-teams neuroimaging replication study, and supervised 5 MSc/MRes/BSc students.

The Alan Turing Institute London, UK

Enrichment Year Fellow & Research Intern Jun 2016 – Sep 2017

- Led an interdisciplinary team of statisticians and computer scientists evaluating classifier-induced neural similarity spaces, resulting in a peer-reviewed publication (*Computational Brain & Behavior*); organised and moderated the institute's PhD Topical Discussions series.

Earlier Research & Public Health Experience (Mexico)

National Institute of Psychiatry Ramón de la Fuente Muñiz (INPRFM) · Social Service Co-Researcher Dec 2010 – Jun 2011

Designed and ran a body-image perception survey ($n = 579$), statistical analysis, and manuscript writing; published in *Rev. Med. IMSS* (2014).

National Institute of Public Health (INSP, CISS) · Research Intern (PROBEI-CCINSHAE Scholarship) Sep 2010 – Mar 2011

Qualitative and ethnographic fieldwork on transnational health practices with returning indigenous migrants in the Mixteca Baja, Oaxaca.

Psychiatric Hospital Fray Bernardino Álvarez · Mental Health Promoter (Volunteer) Jan 2009 – Dec 2009

Community psychiatry outreach, group therapy, and psychodrama workshops in Mexico City.

PUBLICATIONS

Preprints & Industry Articles

Bobadilla-Suarez, S., Suh, B., & Fortin, R. (2026). *Audit the Scaffold, Not the Checkpoint: A Stationarity Dichotomy for Recursive Self-Improvement in Agentic Coding*. arXiv:2609.34924 [cs.LG].

Bobadilla-Suarez, S. (2026). Unity requires physically enforced agreement. *PsyArXiv preprint*. doi:10.31234/osf.io/dkqep_v1 (project page).

Bobadilla-Suarez, S. (2026). When does cortical coherence carry shared content? Compatibility and self-reconstruction as testable conditions. *PsyArXiv preprint*. doi:10.31234/osf.io/fydt3_v1 (project page).

Suh, B., Fortin, R., & **Bobadilla-Suarez, S.** (2026). How "Deep Industry Research Agents" Can Change Your Organization. *Harvard Business Review*, March 2026. hbr.org.

Books & Draft Manuscripts

Bobadilla-Suarez, S. (2025). *Energetically Coherent Computation: A computationally adjacent theory of phenomenal consciousness*. Draft manuscript (v0.5.0, LLM-assisted). doi:10.17605/OSF.IO/42GXN.

Peer-Reviewed Journal Articles

Zhang, X. Y., **Bobadilla-Suarez, S.**, Luo, X., Lemonari, M., Brincat, S. L., Siegel, M., … & Love, B. C. (2025). Adaptive stretching of representations across brain regions and deep learning model layers. *Nature Communications*. doi:10.1038/s41467-025-65231-y.

Bobadilla-Suarez, S., Jones, M., & Love, B. C. (2022). Robust priors for regularized regression. *Cognitive Psychology*, 132, 101444. doi:10.1016/j.cogpsych.2021.101444.

Bobadilla-Suarez, S., Guest, O., & Love, B. C. (2020). Subjective value and decision entropy are jointly encoded by aligned gradients across the human brain. *Communications Biology*, 3(1), 597. doi:10.1038/s42003-020-01315-3.

Botvinik-Nezer, R., Holzmeister, F., Camerer, C., Dreber, A., Huber, J., Johannesson, M., … **Bobadilla-Suarez, S.**, … & Schonberg, T. (2020). Variability in the analysis of a single neuroimaging dataset by many teams. *Nature*, 582(7810), 84–88. doi:10.1038/s41586-020-2314-9.

Bobadilla-Suarez, S., Ahlheim, C., Mehrotra, A., Panos, A., & Love, B. C. (2020). Measures of neural similarity. *Computational Brain & Behavior*, 3(4), 369–383. doi:10.1007/s42113-019-00068-5.

Bobadilla-Suarez, S., & Love, B. C. (2018). Fast or frugal, but not both: Decision heuristics under time pressure. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 44(1), 24–33. doi:10.1037/xlm0000419.

Bobadilla-Suarez, S., Sunstein, C. R., & Sharot, T. (2017). The intrinsic value of control: The propensity to under-delegate in the face of potential gains and losses. *Journal of Risk and Uncertainty*, 54(3), 187–202. doi:10.1007/s11166-017-9259-x.

De Martino, B., **Bobadilla-Suarez, S.**, Noguchi, T., Sharot, T., & Love, B. C. (2017). Social information is integrated into value and confidence judgments according to its reliability. *Journal of Neuroscience*, 37(25), 6066–6074. doi:10.1523/JNEUROSCI.3880-16.2017.

Sunstein, C. R., **Bobadilla-Suarez, S.**, Lazzaro, S., & Sharot, T. (2017). How people update beliefs about climate change: Good news and bad news. *Cornell Law Review*, 102(6).

Bobadilla-Suarez, S., & López-Ávila, A. (2014). Perceived and imagined body image distortion: A possible factor for obesity and overweight in Mexicans [in Spanish]. *Revista Médica del Instituto Mexicano del Seguro Social*, 52(4), 408–414.

Conference Proceedings

Alba Krasovsky, R., **Bobadilla-Suarez, S.**, & Schwarz, D. (2018). Social preference of building materials: Decision-making towards low carbon housing constructions. In *PLEA 2018: Proceedings of the 34th International Conference on Passive and Low Energy Architecture*, Vol. 2 (pp. 778–783).

Bobadilla-Suarez, S., & Love, B. C. (2017). Measures of neural similarity. *Cognitive Computational Neuroscience (CCN 2017)*. ccneuro.org.

Open-Source Code & Public Research Artifacts

OnCorpsAI/audit-the-scaffold (MIT/CC-BY) — Machine-checked Lean 4 proofs, digest-pinned SWE-bench evaluation protocols, and observational multilevel modelling companion to arXiv:2609.34924.

lovelabUCL/dimensionality — Python and MATLAB contributions and release support for estimating functional dimensionality of fMRI data.

INVITED TALKS & INDUSTRY PRESENTATIONS

OnCorps Tech Topics: Agentic coding updates, agentic skills walkthrough, and Agentic Skills Hackathon (4 talks) 2026

OnCorps Tech Topics: "Claude Code (background agents) containerized"; plus OEKS AI coding & DS Journal Club 2025

Psychology Department Seminar, City University (London): "Neural similarity in BOLD response and multi-unit recordings" 2021

CityAI Seminar (City University), Applied ML Group (Forschungszentrum Jülich), WCHN Brain Meeting (UCL), & Neuromatch 3.0: "Neural similarity in BOLD response and multi-unit recordings" 2020

London Judgment and Decision Making Seminar (UCL): "Rethinking heuristics and optimal decision making" 2019

SELECTED CONFERENCE POSTERS

Society for Neuroscience (SfN), Chicago — "The neural link between subjective value and decision entropy" 2019

Interpreting BOLD Symposium, Oxford — "Measures of neural similarity" 2018

Cognitive Computational Neuroscience (CCN), New York & SfN, Washington, D.C. — "Measures of neural similarity" 2017

ESCOP (Paphos) & Workshop on Memory Processes in Judgment and Decision Making (Hölstein) — "Cognitive capacity constraints on heuristic selection" 2015

SUPERVISION & MENTORING

- **Industry Mentoring (OnCorps, 2024 – present):** Direct manager for 2 data scientists; onboarded and mentored data scientists transitioning to product ownership and architectural leadership; trained engineers and non-technical executives on agentic coding workflows.
- **Academic Thesis Supervision (UCL, 2014 – 2021):** Co-supervised 5 students (predominantly MSc Computer Science & Machine Learning) on brain-inspired machine learning projects: Marilena Lemonari (MSc), Nopparada Mingchinda (MRes), Hui Tang (MSc), Jianwen Shen (MSc), and Jay See Tow (BSc).

GRANTS, FELLOWSHIPS & AWARDS

Candidate to National Researcher (Sistema Nacional de Investigadores, SNI)	Mexico
Guarantors of Brain Conference Fee Waiver (∼£250)	UK
The Alan Turing Institute Enrichment Year Stipend (∼£6,000)	2016 – 2017
UCL SLMS Student Conference Fund (∼£500)	UCL
CONACYT Doctoral & Master's Scholarships (over £100,000 combined)	2012 – 2017
PROBEI-CCINSHAE Research Initiation Scholarship, National Institute of Public Health (INSP)	2010 – 2011

CERTIFICATIONS, SERVICE & CONTINUING EDUCATION

- **Databricks Academy (2024):** Completed 15 accreditations and courses covering Lakehouse Fundamentals, AWS Platform Architecture, Data Engineering, Spark Optimisation, ML in Production, Scalable ML, Deep Learning, and LLMs from Foundation Models to Production.
- **Governance & Academic Service:** Member of OnCorps ISO 42001 AIMS Management Board (2025 – present); led 2025 DS Hackathon (Check Writing Agent) and 2026 Agentic Skills Hackathon; organised and moderated The Alan Turing Institute PhD Topical Discussions (2017) and UCL Affective Brain Lab Seminar Series (2014 – 2015).
- **Selected Schools & Courses:** Microsoft Research Summer School in AI (Cambridge, 2017); London Business School Entrepreneurship for Emerging Markets (2015); UCL Enterprise Bootcamp (2015); Plutus Pioneer Program (Haskell, 2021).

SKILLS

LLMs & Agentic AI: Claude Code, LangGraph-style agent harnesses, agent skills, MCP servers/clients, RAG (Qdrant, Databricks Vector Search, sqlite-vec), AWS Bedrock, LiteLLM, OpenAI, Llama/Mistral, Langfuse, MLflow, LLM evaluation (golden sets, deterministic scorers, LLM judges)

ML & Statistics: Robust statistics (MAD), anomaly detection, regularized regression, multilevel/mixed-effects models, sentence-embedding classifiers (SBERT), Bayesian inference, experimental design, PyTorch, scikit-learn

Data & Platform: Databricks (Asset Bundles, Unity Catalog, Lakehouse Federation, Jobs, DLT, SQL), PySpark, Snowflake, Polars, Pandas, PostgreSQL, SQLite, AWS RDS, ElasticSearch

Knowledge Graphs: OWL, RDF/Turtle, SPARQL, Apache Jena Fuseki, FIBO

Software Engineering: Python, TypeScript, R, SQL, Bash, OpenAPI 3.1 contract-first design, semver gates, SDK codegen, GitHub Actions CI, Docker, Kubernetes, FastAPI, Dash, R Shiny, mutation testing, Lean 4 (formal proofs)

AI Governance & Leadership: ISO 42001 AIMS, ISMS tool reviews, Red Zone human-in-the-loop controls, fail-closed design, programme leadership, ELT/board reporting, hiring & mentoring

Computational Neuroscience: fMRI MVPA & representational similarity analysis (RSA), model-based neuroimaging (SPM, FSL), eye tracking (EyeLink II), MATLAB, HPC clusters

Languages: English (fluent), Spanish (native); basic familiarity with Italian, French, and Portuguese