



From “Agents” to Autonomy: A Practical Framework for Agentic AI (Levels 1-5) (<https://opendatascience.com/from-agents-to-autonomy-a-practical-framework-for-agentic-ai-levels-1-5/>)

MODELING (<https://opendatascience.com/category/modeling/>) AGENTIC AI (<https://opendatascience.com/tag/agentic-ai/>) AI AGENTS (<https://opendatascience.com/tag/ai-agents/>) posted by ODSC Community (<https://opendatascience.com/user/odsc-community/>) 📅 December 3, 2025

🔖 Agentic AI 8 (<https://opendatascience.com/tag/agentic-ai/>) AI Agents 5 (<https://opendatascience.com/tag/ai-agents/>)

Business leaders today are inundated with claims that **artificial general intelligence (AGI)** – a human-level, autonomous AI – is just around the corner. Headlines trumpet looming “**super-agents**” (<https://www.axios.com/2025/01/19/ai-superagent-openai-meta#:~:text=Architects%20of%20the%20leading%20generative,to%20do%20complex%20human%20tasks>) that might soon replace knowledge workers, stoking both excitement and anxiety

Agentic AI has become the enterprise buzzword of 2025, referring to AI systems that can act with some degree of autonomy to achieve goals, *emphasis on ‘degree of autonomy’*. The Society of Automotive Engineers addressed a similar challenge with the advent of self-driving cars, defining expected functionality and the necessary human oversight for safe and effective operation.

Just as autonomous vehicles use SAE Levels 0–5 (https://www.sae.org/standards/j3016_202104-taxonomy-definitions-terms-related-driving-automation-systems-road-motor-vehicles), AI systems need a shared vocabulary for autonomy. This post turns a live, complicated discussion into a clear, practical guide you can use to plan, scope, and safely deploy agentic systems inside real organizations.

Get any questions? I'm happy to help.



Defining Autonomy

		SAE LEVEL 0™	SAE LEVEL 1™	SAE LEVEL 2™	SAE LEVEL 3™	SAE LEVEL 4™	SAE LEVEL 5™
What does the human in the driver's seat have to do?		You are driving whenever these driver support features are engaged – even if your feet are off the pedals and you are not steering			You are not driving when these automated driving features are engaged – even if you are seated in “the driver's seat”		
		You must constantly supervise these support features; you must steer, brake or accelerate as needed to maintain safety			When the feature requests, you must drive	These automated driving features will not require you to take over driving	
Copyright © 2021 SAE International.							
		These are driver support features			These are automated driving features		
What do these features do?		These features are limited to providing warnings and momentary assistance	These features provide steering OR brake/acceleration support to the driver	These features provide steering AND brake/acceleration support to the driver	These features can drive the vehicle under limited conditions and will not operate unless all required conditions are met		This feature can drive the vehicle under all conditions
	Example Features	• automatic emergency braking • blind spot warning • lane departure warning	• lane centering OR • adaptive cruise control	• lane centering AND • adaptive cruise control at the same time	• traffic jam chauffeur	• local driverless taxi • pedals/steering wheel may or may not be installed	• same as level 4, but feature can drive everywhere in all conditions

(<https://opendatascience.com/wp-content/uploads/2025/12/IC1.png>)

(SAE Levels of Driving Automation (<https://www.sae.org/news/blog/sae-levels-driving-automation-clarity-refinements>), Updated May 2021)

It's worth noting that the functionality gains from one level to the next in the case of self-driving cars are non-linear. In fact, it's been nearly a decade since Elon first promised that fully self-driving cars are around the corner, yet Teslas (no matter how great they operate *most of the time*) still require humans to be able to seize control at any moment. And yes, we have Waymo operating in *limited* environments

The problems to solve (and the cost to solve them) jump exponentially at each level, and we posit that, just like fully autonomous self-driving cars at level five that can operate anywhere, we are a far way off from seeing level five agentic AI operating at the “PhD” level as Sam Altman describes it. And the authors both firmly believe that LLMs are not the architecture to get us to level five, but that is a discussion for a separate post.



Register for ODSC AI East 2026 - Now 70% off!

(https://www.eventbrite.com/e/odsc-ai-east-2026-conference-open-data-science-conference-tickets-1780200771429?aff=odscailblog&_eboga=false)



(https://www.eventbrite.com/e/odsc-ai-east-2026-conference-open-data-science-conference-tickets-1780200771429?aff=odscailblog&_eboga=false)In-person conference | April 28th-30th, 2026 | Boston, MA

Meet leading experts, upskill with hands-on workshops, and connect with thousands of data science and AI practitioners shaping the next wave of innovation.

Save an additional 10% with this link! (https://www.eventbrite.com/e/odsc-ai-east-2026-conference-open-data-science-conference-tickets-1780200771429?aff=odscailblog&_eboga=false)

The Five Levels of Agentic AI

The challenge with the current narrative around Agentic AI is that it paints all autonomy with the same brush, with many claims of models “quickly improving” to “superhuman levels.” These arguments miss the point. Autonomy isn't a binary “yes” or “no.” It occurs along a continuum of generalization, human-AI collaboration, and oversight.

To make the phrase “Agentic AI” more productive, we propose the taxonomy below, inspired by the National Highway Traffic Safety Administration's definitions of self-driving. 

	Level 1 — Deterministic Task Bot	Level 2 — Preparatory Agent	Level 3 — Narrow Operator	Level 4 — Semi-Autonomous Specialist	Level 5 — Autonomous Problem Solver
What it does	Executes a narrow, predefined action	- Drafts, proposes, or partially executes multi-step tasks Always requires review	- Handles routine workflows end-to-end - Escalates uncertain cases	Operates correctly the vast majority of the time (think ~“98%”), including planning, tool use, and recovery from common failures	- Pursues open-ended goals - Decomposes novel problems - Generates new knowledge or strategies (the “PhD-level” analogy)**
Example	- Password reset workflow - Triaging tickets - Filing a form - Scheduling a meeting with fixed templates	- Draft contract redlines - Compile weekly business summary - Propose a JIRA triage plan	- Close-out support tickets - Run weekly reporting - Push low-risk config changes - Prepare e-commerce listings	- Price monitoring and reordering with budgets - Lead-to-opportunity routing - L2/L3 IT ops playbooks with rollback - Travel booking within strict constraints	- New algorithmic ideas - Novel drug hypotheses - Independent market entry strategies
Human oversight	- Required for edge cases - Output is easy to verify	Mandatory review/approval before anything hits production or customers. Treat like a new hire you’re training	Spot checks; “trust, but verify.” Humans remain on the hook for final accountability	- Humans on the loop - Approvals only for out-of-policy events, anomalies, or high-impact decisions	Policy-level supervision and after-action audits—not step-by-step review
Tooling/permissions	- Single system or API - No spending authority	- Multiple tools (email, docs, trackers) in read/limited write No autonomous spending	- Broader tool access - Constrained writes - May trigger pre-approved micro-actions (e.g., issue credits up to \$X)	- Multi-tool orchestration Bounded time/budget authority - Strong guardrails (rate limits, RBAC, circuit breakers)	- Broad read/write with explicit scoping - Significant resources - Programmatic ethics/compliance gates
Where it shines	High-volume, low-variance tasks you could have automated even pre-GenAI—now easier to build and maintain	Knowledge work with templates, checklists, and strong reviewer SLAs	Stable processes with known exceptions and good audit trails	Medium-risk domains where the cost of delay > cost of occasional correction, and where governance is robust	Technically exciting, but unnecessary— and unsafe—for most enterprise workflows today. Treat as research, not production

**Note: The authors do not believe that this level of autonomy is reachable in the short term.

Level 1 — Deterministic Task Bot

What it does: Executes a narrow, predefined action with minimal reasoning.

Example: Password reset workflow; triaging tickets; filing a form; scheduling a meeting with fixed templates.

Human oversight: Required for edge cases; output is easy to verify.

Tooling/permissions: Single system or API; no spending authority.

Where it shines: High-volume, low-variance tasks you could have automated even pre-GenAI—now easier to build and maintain.

Level 2 — Preparatory Agent

What it does: Drafts, proposes, or partially executes multi-step tasks; **always** requires review.

Example: Draft contract redlines; compile weekly business summary; propose a JIRA triage plan.

Human oversight: Mandatory review/approval before anything hits production or customers. Treat like a new hire you’re training.

Tooling/permissions: Multiple tools (email, docs, trackers) in read/limited write; **no** autonomous spending.

Where it shines: Knowledge work with templates, checklists, and strong reviewer SLAs.



Level 3 — Narrow Operator

What it does: Handles routine workflows end-to-end; escalates uncertain cases.
Example: Close-out support tickets, run weekly reporting, push low-risk config changes, prepare e-commerce listings.
Human oversight: Spot checks; “trust, but verify.” Humans remain on the hook for final accountability.
Tooling/permissions: Broader tool access; constrained writes; may trigger **pre-approved** micro-actions (e.g., issue credits up to \$X).
Where it shines: Stable processes with known exceptions and good audit trails.

Level 4 — Semi-Autonomous Specialist

What it does: Operates correctly the vast majority of the time (think ~98%), including planning, tool use, and recovery from common failures.
Example: Price monitoring and reordering with budgets; lead-to-opportunity routing; L2/L3 IT ops playbooks with rollback; travel booking **within** strict constraints.
Human oversight: Humans in the loop; approvals only for out-of-policy events, anomalies, or high-impact decisions.
Tooling/permissions: Multi-tool orchestration; **bounded** time/budget authority; strong guardrails (rate limits, RBAC, circuit breakers).
Where it shines: Medium-risk domains where the cost of delay > cost of occasional correction, and where governance is robust.

Level 5 — Autonomous Problem Solver

What it does: Pursues open-ended goals, decomposes novel problems, and generates new knowledge or strategies (the “PhD-level” analogy).
Example: New algorithmic ideas; novel drug hypotheses; independent market entry strategies.
Human oversight: Policy-level supervision and after-action audits—**not** step-by-step review.
Tooling/permissions: Broad read/write with explicit scoping; significant resources; programmatic ethics/compliance gates.
Reality check: Technically exciting, but unnecessary—and unsafe—for most enterprise workflows today. Treat as research, not production.



Build, Deploy & Scale Autonomous AI Agents in a 3-Week Hands-on Virtual Summit



AGENTIC AI SUMMIT

Master Agentic AI

SUPER EARLY BIRD PASS AVAILABLE NOW

3-Week Hands-On Training

Virtual | JAN 21-FEB 5, 2026

([https://www.summit.ai/?](https://www.summit.ai/?__hstc=39712252.127a4914be2745389e3be98e05364b99.1749618174162.1763616159107.1764825017850.14&__hssc=39712252.1.1764825017850&__hsfp=331823229)

[__hstc=39712252.127a4914be2745389e3be98e05364b99.1749618174162.1763616159107.1764825017850.14&__hssc=39712252.1.1764825017850&__hsfp=331823229](https://www.summit.ai/?__hstc=39712252.127a4914be2745389e3be98e05364b99.1749618174162.1763616159107.1764825017850.14&__hssc=39712252.1.1764825017850&__hsfp=331823229))Attend the Agentic AI Summit, a 3-week virtual, hands-on program designed to help you build, deploy, and scale autonomous AI agents.

Jan 21 – Feb 5, 2026 | Virtual

👉 Register now with Super Early Bird Discount: [https://www.summit.ai/](https://www.summit.ai/?__hstc=39712252.127a4914be2745389e3be98e05364b99.1749618174162.1763616159107.1764825017850.14&__hssc=39712252.1.1764825017850&__hsfp=331823229) ([https://www.summit.ai/?](https://www.summit.ai/?__hstc=39712252.127a4914be2745389e3be98e05364b99.1749618174162.1763616159107.1764825017850.14&__hssc=39712252.1.1764825017850&__hsfp=331823229)

[__hstc=39712252.127a4914be2745389e3be98e05364b99.1749618174162.1763616159107.1764825017850.14&__hssc=39712252.1.1764825017850&__hsfp=331823229](https://www.summit.ai/?__hstc=39712252.127a4914be2745389e3be98e05364b99.1749618174162.1763616159107.1764825017850.14&__hssc=39712252.1.1764825017850&__hsfp=331823229))

Or get complimentary access with any paid ODSC AI East 2026 Pass! (<https://www.eventbrite.com/e/odsc-ai-east-2026-conference-open-data-science-conference-tickets-1780200771429?aff=odsciblog&eboga=false>)

Learn more & sign up here ([https://www.summit.ai/?](https://www.summit.ai/?__hstc=39712252.127a4914be2745389e3be98e05364b99.1749618174162.1763616159107.1764825017850.14&__hssc=39712252.1.1764825017850&__hsfp=331823229)

[__hstc=39712252.127a4914be2745389e3be98e05364b99.1749618174162.1763616159107.1764825017850.14&__hssc=39712252.1.1764825017850&__hsfp=331823229](https://www.summit.ai/?__hstc=39712252.127a4914be2745389e3be98e05364b99.1749618174162.1763616159107.1764825017850.14&__hssc=39712252.1.1764825017850&__hsfp=331823229))

One Size of Autonomy Does Not Fit All

Our taxonomy is primarily organized around user control and tool scope as these are the most visible features to end users. That said, we recognize that other dimensions may be helpful for classifying the autonomy of agents at different levels.



- 1. **Task complexity:** The amount of human time typically required to perform a task is a reasonable proxy for complexity. One-shot actions (e.g., sending an email) are more straightforward than complex, multi-step tasks (e.g., researching various options before purchasing).
- 2. **ROI potential:** A very direct measure of agents is the potential for economic impact, cost savings, or revenue generation. OpenAI recently introduced a benchmark, GDPVal (<https://openai.com/index/gdpval/>), that evaluates various AI models on real-world economic impacts.
- 3. **Domain risk & consequence:** Legal, healthcare, finance, and safety-critical ops require tighter guardrails than marketing or internal ops. Such industries will have different tolerance levels for the adoption of autonomous agents. Auditability and explainability may be required as an additional layer of implementation here.
- 4. **Multi-modality:** The majority of agents in production today work well only with plain-text documents, with strict limitations for reading images, tables, charts, audio, and video. As models and other engineering tools are implemented, agents will be able to consume more data sources to make more informed decisions.

Making Agentic AI More Productive with a Taxonomy

A shared taxonomy gives us a cleaner way to talk about autonomy. Right now, everything from a scripted workflow to a multi-tool planner gets labeled an “agent,” which blurs important differences in oversight, permissions, and expected behavior.

Most organizations will harvest outsized gains from well-governed Level 2–4 systems long before Level 5 is necessary. If you take one next step this quarter, pilot a single Level-2 workflow, define necessary human oversight and success metrics, and iterate toward bounded autonomy. The result isn't just faster execution; it's trustworthy automation that your teams understand, your auditors can verify, and your customers actually benefit from, at a lower total cost of ownership.

We encourage you to treat this framework for agentic AI as a living document. Our goal is a shared language that makes it easier to argue less about what an “agent” is and focus more on what type of autonomy is needed. If you are mapping your own workflows to these levels, or finding edge cases that do not fit, we would like to hear from you.

As scientific and engineering breakthroughs accelerate, we expect the average level of these agents to rise alongside them. Regardless of the pace of change, it is clear to us that setting clear user expectations is what will drive successful adoption.

About the Authors

Cal Al-Dhubaib (<https://www.linkedin.com/in/dhubaib/>) is a data scientist, entrepreneur, and innovator in responsible artificial intelligence, specializing in high-risk sectors such as healthcare, energy, and defense. He is the Head of AI and Data Science at Further. Cal frequently speaks on topics including AI ethics, change management, data literacy, and the unique challenges of implementing AI solutions in high-risk industries. His insights have been featured in numerous publications such as Forbes, Ohio, the Marketing AI Institute, Open Data Science, and AI Business News. Cal has also received recognition among Crain's Cleveland Notable Immigrant Leaders, Notable Entrepreneurs, and most recently, Notable Technology Executives.

Ivan Lee (<https://www.linkedin.com/in/iylee/>) graduated with a Computer Science B.S. from Stanford University, then dropped out of his master's degree to found his first mobile gaming company Loki Studios. After raising institutional funding and building a profitable game, Loki was acquired by Yahoo. Lee spent the next 10 years building AI products at Yahoo and Apple and discovered there was a gap in serving the rapid evolution of Natural Language Processing (NLP) technologies. He built Datasaur to focus on democratizing access to NLP and LLMs. Datasaur has raised \$8m in venture funding from top-tier investors such as Initialized Capital, Greg Brockman (President, OpenAI) and Calvin French-Owen (CTO, Segment) and serves companies such as Google, Netflix, Qualtrics, Spotify, the FBI and more.

0

Twitter

0

Pin

ABOUT AUTHOR



ODSC Community (<https://opendatascience.com/user/odsc-community/>)
The Open Data Science community is passionate and diverse, and we always welcome contributions from data science professionals! All of the articles under this profile are from our community, with individual authors mentioned in the text itself.





Stay connected! Get curated newsletters every week

First Name*

Seamas

Last name*

Email*

seamas.yau@gmail.com

Country/Region

Please Select

From time to time, we'd like to contact you with other related content and offers. You may unsubscribe at any time. [Privacy Policy \(https://opendatascience.com/odsc-privacy-policy/\)](https://opendatascience.com/odsc-privacy-policy/)

☐

I agree to receive communications from ODSC.

Subscribe

LATEST POSTS

View all 



(<https://opendatascience.com/from-agents-to-autonomy-a-practical-framework-for-agentic-ai-levels-1-5/>)

From “Agents” to Autonomy: A Practical Framework for Agentic AI (Levels 1–5) (<https://opendatascience.com/from-agents-to-autonomy-a-practical-framework-for-agentic-ai-levels-1-5/>)

posted by [ODSC COMMUNITY \(https://opendatascience.com/user/odsc-community/\)](https://opendatascience.com/user/odsc-community/)

🕒 Dec 3, 2025

Business leaders today are inundated with claims that artificial general intelligence (AGI) – a human-level, autonomous...



(<https://opendatascience.com/oracle-initiated-at-overweight-as-wells-fargo-sees-room-to-run-in-ai-super-cycle/>)

Oracle Initiated at Overweight as Wells Fargo Sees “Room to Run” in AI Super-Cycle

(<https://opendatascience.com/oracle-initiated-at-overweight-as-wells-fargo-sees-room-to-run-in-ai-super-cycle/>)

posted by [ODSC TEAM \(https://opendatascience.com/user/admin_odsc/\)](https://opendatascience.com/user/admin_odsc/)

🕒 Dec 3, 2025

Wells Fargo initiated coverage of Oracle with an Overweight rating and a \$280 price target, citing...



(<https://opendatascience.com/the-ai-frenzy-is-fueling-a-new-global-memory-chip-crisis/>)

The AI Frenzy is Fueling a New Global Memory Chip Crisis (<https://opendatascience.com/the-ai-frenzy-is-fueling-a-new-global-memory-chip-crisis/>)

posted by [ODSC TEAM \(https://opendatascience.com/user/admin_odsc/\)](https://opendatascience.com/user/admin_odsc/)

🕒 Dec 3, 2025

A surge in AI investment is straining the world's memory-chip supply chain, creating shortages that now...





POPULAR BLOGS



(<https://opendatascience.com/15-open-healthcare-datasets-2024-update/>) **15 Open Healthcare Datasets – 2024 Update**
(<https://opendatascience.com/15-open-healthcare-datasets-2024-update/>)



(<https://opendatascience.com/top-10-open-source-ai-agent-frameworks-to-know-in-2025/>) **Top 10 Open-Source AI Agent Frameworks to Know in 2025**
(<https://opendatascience.com/top-10-open-source-ai-agent-frameworks-to-know-in-2025/>)



(<https://opendatascience.com/the-ai-frenzy-is-fueling-a-new-global-memory-chip-crisis/>) **The AI Frenzy is Fueling a New Global Memory Chip Crisis**
(<https://opendatascience.com/the-ai-frenzy-is-fueling-a-new-global-memory-chip-crisis/>)



(<https://opendatascience.com/from-agents-to-autonomy-a-practical-framework-for-agentic-ai-levels-1-5/>) **From “Agents” to Autonomy: A Practical Framework for Agentic AI (Levels 1–5)**
(<https://opendatascience.com/from-agents-to-autonomy-a-practical-framework-for-agentic-ai-levels-1-5/>)



(<https://opendatascience.com/3-examples-of-ai-helping-police-solve-cold-cases/>) **3 Examples of AI Helping Police Solve Cold Cases**
(<https://opendatascience.com/3-examples-of-ai-helping-police-solve-cold-cases/>)



TAGS

Machine Learning	111	(https://opendatascience.com/tag/machine-learning/)	Podcast	87	(https://opendatascience.com/tag/podcast/)
Azure	70	(https://opendatascience.com/tag/azure/)	East 2022	65	(https://opendatascience.com/tag/east-2022/)
ODSC East 2015 Speaker Slides	64	(https://opendatascience.com/tag/odsc-east-2015speaker-slides/)	East 2024	58	(https://opendatascience.com/tag/east-2024/)
East 2025	58	(https://opendatascience.com/tag/east-2025/)	West 2024	52	(https://opendatascience.com/tag/west-2024/)
Microsoft	50	(https://opendatascience.com/tag/microsoft/)	West 2025	50	(https://opendatascience.com/tag/west-2025/)
West 2022	49	(https://opendatascience.com/tag/west-2022/)	Deep Learning	48	(https://opendatascience.com/tag/deep-learning/)
East 2020	48	(https://opendatascience.com/tag/east-2020/)	East 2023	47	(https://opendatascience.com/tag/east-2023/)
West 2021	46	(https://opendatascience.com/tag/west-2021/)	Accelerate AI	43	(https://opendatascience.com/tag/accelerate-ai/)



cybersecurity	42 (https://opendatascience.com/tag/cybersecurity/)	East 2021	42 (https://opendatascience.com/tag/east-2021/)
Conferences	41 (https://opendatascience.com/tag/conferences/)	West 2023	40 (https://opendatascience.com/tag/west-2023/)

RELATED POSTS



ABOUT US

Proactively envisioned multimedia based expertise and cross-media growth strategies. Seamlessly visualize quality intellectual capital without superior collaboration and idea-sharing. Holistically pontificate installed base portals after maintainable products.

CONTENT MAP

Modeling (https://opendatascience.com/category/modeling/)	1845	AI and Data Science News (https://opendatascience.com/category/news/)	1078
Conferences (https://opendatascience.com/category/conferences/)	836	Featured Post (https://opendatascience.com/category/featured-post/)	827
Business + Management (https://opendatascience.com/category/accelerate-ai/)	560	Tools & Languages (https://opendatascience.com/category/tools-languages/)	365
Machine Learning (https://opendatascience.com/category/modeling/machine-learning/)	353	NLP & LLMs (https://opendatascience.com/category/modeling/nlp-text-analytics/)	318
Deep Learning (https://opendatascience.com/category/modeling/deep-learning/)	233	Career Insights (https://opendatascience.com/category/careers/)	223
Blogs from ODSC Speakers (https://opendatascience.com/category/conferences/odsc-speaker/)	202	Blog (https://opendatascience.com/category/blog/)	201
Data Visualization (https://opendatascience.com/category/modeling/data-visualization/)	137	Statistics (https://opendatascience.com/category/modeling/statistics/)	126
R (https://opendatascience.com/category/tools-languages/r/)	124	Python (https://opendatascience.com/category/tools-languages/python/)	115
Predictive Analytics (https://opendatascience.com/category/modeling/predictive-analytics/)	99	Technology (https://opendatascience.com/category/accelerate-ai/technology/)	77
Guest contributor (https://opendatascience.com/category/guest-contributor/)	74	Research (https://opendatascience.com/category/modeling/research-modeling/)	72

TAGS

Machine Learning	111 (https://opendatascience.com/tag/machine-learning/)	Podcast	87 (https://opendatascience.com/tag/podcast/)
Azure	70 (https://opendatascience.com/tag/azure/)	East 2022	65 (https://opendatascience.com/tag/east-2022/)
ODSC East 2015 Speaker Slides	64 (https://opendatascience.com/tag/odsc-east-2015speaker-slides/)	East 2024	58 (https://opendatascience.com/tag/east-2024/)
East 2025	58 (https://opendatascience.com/tag/east-2025/)	West 2024	52 (https://opendatascience.com/tag/west-2024/)
Microsoft	50 (https://opendatascience.com/tag/microsoft/)	West 2025	50 (https://opendatascience.com/tag/west-2025/)
West 2022	49 (https://opendatascience.com/tag/west-2022/)	Deep Learning	48 (https://opendatascience.com/tag/deep-learning/)
East 2020	48 (https://opendatascience.com/tag/east-2020/)	East 2023	47 (https://opendatascience.com/tag/east-2023/)



West 2021	46 (https://opendatascience.com/tag/west-2021/)	Accelerate AI	43 (https://opendatascience.com/tag/accelerate-ai/)
cybersecurity	42 (https://opendatascience.com/tag/cybersecurity/)	East 2021	42 (https://opendatascience.com/tag/east-2021/)
Conferences	41 (https://opendatascience.com/tag/conferences/)	West 2023	40 (https://opendatascience.com/tag/west-2023/)

ODSC PRIVACY POLICY

View ODSC Privacy Policy (<https://opendatascience.com/odsc-privacy-policy/>)

