

AI Skills in Pharma Job Postings: Roles & Training Needs

IntuitionLabs.ai · Adrien Laurent · 2026-09-05 · ai skills pharma · pharma job postings · machine learning jobs pharma · life sciences ai talent · data science pharmaceutical industry · ai training pharma · ai roles drug development · pharma workforce ai · artificial intelligence hiring life sciences



Executive Summary

Artificial intelligence (AI) keywords remain rare in pharmaceutical job postings by every large-sample measure available as of September 2026, yet the roles that do carry them ask for unusually specific technical skills. Lightcast's life-sciences workforce analysis found just over 2% of UK Life Sciences job postings between April 2023 and March 2024 mentioned AI-related skills, versus less than 1% across the broader UK labor market (www.datocms-assets.com), and PwC's 2026 Global AI Jobs Barometer separately placed the health sector at the lowest end of AI-skill job-posting share among tracked industries, at under 1% (www.pwc.com). Yet this report's direct sample of fifteen live **AI and machine-learning job postings**, fetched from six major drugmakers' career sites in late August and early September 2026, found named requirements spanning Python and PyTorch, large language models (LLMs) and retrieval-augmented generation (RAG), reinforcement learning from human feedback (RLHF), and cheminformatics libraries such as RDKit and DeepChem, concentrated in research and development (R&D), platform-engineering, and computational-chemistry roles.

Occupational data situates this demand within faster-than-average growth: the US Bureau of Labor Statistics (BLS) projects Data Scientist employment to grow 35% from 2025 to 2035 (www.bls.gov), with a median annual wage of \$120,230 in May 2025 (www.bls.gov). PwC found AI-skill-specific postings growing roughly eight times faster than the total jobs market (69% versus 9%) (www.pwc.com), and Indeed's Hiring Lab counted 822 distinct "AI-touched" US job titles by early 2026, up from 264 in early 2022 (hiringlab.indeed.com).

The training ecosystem built to close this gap operates on three levels. Enterprise-wide accreditation programs, such as AstraZeneca's tiered initiative that grew from over 12,000 certified employees in April 2025 (www.astrazeneca.com) to over 37,700 by October 2025 (www.astrazeneca.com), target broad workforce literacy. Professional-society courses, including DIA's AI-in-pharmacovigilance training (www.diaglobal.org) and RAPS's AI in Regulatory Conference (live.raps.org), target regulated functions specifically. New university degree programs, led by UCSF's Master of Science in AI and Computational Drug Discovery and Development, described as the first of its kind in the United States (pharmacy.ucsf.edu), aim to graduate AI-literate scientists directly.

The gap these pipelines address is well documented: a 2019 Tufts Center for the Study of Drug Development (CSDD) and Drug Information Association (DIA) survey found staff skills was the top-cited barrier to **AI implementation in drug development**, at 55% (www.ebi.ac.uk), and PhRMA's more recent workforce analysis projects the US biopharmaceutical industry will need to hire or reskill roughly 30,000 workers annually through 2034 (phrma.org). Read together, the evidence supports treating "AI skills" in pharma hiring as a defined, classifiable set of named tools and methods tied to specific roles, methodology, and training pathway, documented in full in the sections that follow, rather than as a single generic hiring requirement.

Introduction and Background

Pharmaceutical employers spent 2025 and 2026 rewriting job descriptions around artificial intelligence (AI), but how much rewriting, and in which roles, is easier to assert than to measure with a defined method. This report analyzes what pharmaceutical and life-sciences employers actually write into current AI-related job postings as of September 5, 2026: which specific technical skills appear (Python, PyTorch, large language models (LLMs), cheminformatics libraries), which job titles carry them, and what training pathways exist to close the resulting gap. It treats "AI skills" narrowly and defines the classification schema used, rather than treating the term as a synonym for general digital transformation.

The starting point is a tension in the published labor-market data. Lightcast's life-sciences workforce study found that just over 2% of UK Life Sciences job postings between April 2023 and March 2024 mentioned AI-related skills, compared with less than 1% across the broader UK labour market (www.datocms-assets.com). PwC's 2026 Global AI Jobs Barometer independently places health-sector postings at the low end of AI-skill penetration among the industries it tracks, at less than 1% (www.pwc.com). At the same time, Indeed's Hiring Lab counted 822 distinct "AI-touched" job titles across the US labor market by the first quarter of 2026, up from 264 in the first quarter of 2022 (hiringlab.indeed.com), and a 2019 Tufts Center for the Study of Drug Development (CSDD) and Drug Information Association (DIA) survey found that inadequate staff skill sets were already a significant hurdle to **AI adoption in biopharmaceutical development** (www.diaglobal.org). Both things are true simultaneously: AI-specific keywords remain rare as a share of all pharmaceutical postings, while the roles that do carry them, and the difficulty of filling those roles, have grown substantially.

This report does not attempt to re-run a company-level hiring census; a companion analysis already covers named drugmakers' AI organization-building, layoffs, and compensation signals in detail (intuitionlabs.ai). Nor does it re-litigate the widely cited pharma AI skills-gap survey statistics; a separate companion report covers those reskilling-economics figures (intuitionlabs.ai). Instead, this report is organized around a reproducible question: what does the word "AI" actually mean inside a pharmaceutical job posting, in which roles does it appear, and what training closes the gap between the two? IntuitionLabs approaches this question from an implementation perspective: **"Develop role-specific capability across regulatory, medical writing, clinical operations, medical affairs, quality, CMC, and commercial teams"** (intuitionlabs.ai) is a materially different exercise from measuring what employers already ask for in writing, which is this report's focus.

Methodology: Classifying AI Skills in Pharma Job Postings

Three complementary evidence sources were combined, each with a stated method, so a reader can reproduce or extend the analysis.

Published posting-level taxonomies. Lightcast, a labor-market analytics firm, classified roughly 200 skills, knowledge items, and abilities into AI-related clusters such as machine learning, robotics, and natural language processing (NLP) across its life-sciences postings database (www.datocms-assets.com). Indeed's Hiring Lab, separately, applied a list of 591 AI- and generative-AI-related keywords, spanning concepts such as deep learning, GenAI, agentic AI, and NLP, to classify US job postings between July 2024 and June 2025 (hiringlab.indeed.com). Neither firm publishes a pharma-only cut of every metric it reports, so this report is explicit about which figures are life-sciences-specific versus economy-wide.

A direct sample of current employer postings. To ground the taxonomy in what pharmaceutical employers write today, this report's research drew a non-random, illustrative sample of fifteen live AI- and machine-learning-related job postings, fetched directly from the official career sites of six large drugmakers (Novartis, Bristol Myers Squibb, Merck & Co. (known as MSD outside the United States and Canada), Amgen, Pfizer, and Johnson & Johnson) between late August and early September 2026. Each posting's exact required-skill language was extracted and classified into the categories used in the next section. This sample is a snapshot, not a census: postings rotate on a weeks-to-months cycle, sample selection favored companies with easily indexed career pages, and a search for equivalent live postings at AstraZeneca, Eli Lilly, Genentech/Roche, GSK, and Sanofi did not return a fetchable, unexpired posting during the same window, a limitation discussed further below.

Occupational and skills-registry data. The US Bureau of Labor Statistics (BLS) Occupational Outlook Handbook and the US Department of Labor's O*NET OnLine database provide standardized, government-maintained definitions of the occupations (Data Scientists, Computer and Information Research Scientists, Statisticians, Biochemists and Biophysicists) into which most pharma AI postings fall, including which specific software tools each occupation's profile flags as a "Hot Technology," O*NET's designation for tools with fast-growing employer demand.

Classification schema. For this report, an "AI skill" is any explicitly named machine-learning method, model family, framework, or tool (for example, PyTorch, TensorFlow, LLMs, retrieval-augmented generation (RAG), reinforcement learning, or a domain-specific cheminformatics library), rather than the bare word "AI." This mirrors Indeed's own finding that generic references dominate: nearly 74% of AI-related postings it analyzed used the general term "AI,"

while only 2% named a specific product such as ChatGPT ([hiringlab.indeed.com](https://www.hiringlab.indeed.com)). A posting-level count of "AI mentions" therefore overstates concrete technical requirements unless, as this report does, it separately tracks named tools and methods.

Limitations. The direct-sample method above is a reproducible illustration, not a statistically representative census: fifteen postings across six employers cannot support population-level percentage estimates the way Lightcast's or Indeed's much larger indexed databases can. Job postings also do not necessarily reflect actual hiring outcomes, required skills are sometimes aspirational or copied from templates, and posting language varies by recruiter. Where this report cites a percentage, it attributes it to the specific database (UK Life Sciences postings, US economy-wide postings) that produced it, and does not extrapolate a UK or economy-wide figure onto the US pharmaceutical sector specifically.

AI Skill Categories in Current Pharma Job Postings

Table 1 below summarizes the specific AI-related skills, tools, and frameworks found in the fifteen live postings sampled directly from drugmaker career sites, grouped into the categories this report uses throughout.

Table 1: AI Skill Categories Observed in a Sample of Current Pharma Employer Job Postings (accessed late August to early September 2026)

Skill Category	Specific Skills or Tools Named	Example Employer and Role
Core machine learning and deep learning frameworks	Python, PyTorch, TensorFlow, Spark, Git version control	Novartis Director and Group Head, Applied AI, requiring "strong coding proficiency in Python and deep learning frameworks, with experience using version control systems such as Git" (novartis.wd3.myworkdayjobs.com); Novartis Director, Data Science, requiring proficiency in "Python, Spark, TensorFlow, and PyTorch" (novartis.wd3.myworkdayjobs.com)
Generative AI, LLMs, and agentic systems	LLMs, retrieval-augmented generation (RAG), agentic workflow design, foundation models, reasoning models	Bristol Myers Squibb Manager, Pharmaceutical Product Development GenAI & Data Science, naming "Large Language Models, Reasoning Models, Retrieval-Augmented Generation (RAG), Agentic Workflow Design" (bristolmyerssquibb.wd5.myworkdayjobs.com); Pfizer Translational AI Engineer, covering "generative AI, agentic workflows, predictive models, foundation models, retrieval-augmented systems, and fine-tuned model architectures" (pfizer.wd1.myworkdayjobs.com)
Domain-specific scientific AI (cheminformatics, imaging, structural biology)	RDKit, OpenEye, DeepChem, SMILES/SMARTS notations, MONAI, Hugging Face, Lightning	Pfizer Senior ML/AI Cheminformatics Scientist, requiring "cheminformatics libraries and topics (RDKit, OpenEye, DeepChem, InChI, SMILES/SMARTS)" (pfizer.wd1.myworkdayjobs.com); Novartis Data Scientist, AI for Biomedical Imaging, requiring the "scientific deep learning stack (e.g., PyTorch, Hugging Face, Lightning, MONAI)" (novartis.wd3.myworkdayjobs.com)

Skill Category	Specific Skills or Tools Named	Example Employer and Role
Reinforcement learning and advanced model training	Deep reinforcement learning (RL), policy optimization, reinforcement learning from human feedback (RLHF)	Amgen Principal Machine Learning Engineer / AI Platform Lead, requiring "expertise in Reinforcement Learning (Deep RL, Policy Optimization, RLHF)" (amgen.wd1.myworkdayjobs.com)
MLOps, cloud, and deployment engineering	AWS, MLFlow, Docker, GitHub Actions continuous integration/continuous delivery (CI/CD)	Merck (MSD) Associate Director, AI/ML Engineering, requiring "Python, Torch, AWS, MLFlow, Docker, GitHub, and GitHub Actions CI/CD workflows" (msd.wd5.myworkdayjobs.com); Bristol Myers Squibb Senior Specialist, AI Engineering & Enablement, covering "deployment and operation of AI solutions on AWS cloud services and enterprise AI platforms" (bristolmyerssquibb.wd5.myworkdayjobs.com)
General data science and statistical programming	Python, R, SQL	Johnson & Johnson Senior Data Scientist, Service Innovation, requiring "advanced proficiency in Python and SQL for data processing, statistical analysis, machine learning, and query optimization" (www.careers.jnj.com); Amgen Scientist, Digital Discovery, preferring "Proficiency in Python and/or R" (amgen.wd1.myworkdayjobs.com)

Across this sample, no two postings used identical skill lists, but three patterns recur. First, Python is a near-universal baseline: it appears, alone or alongside a specific framework, in the great majority of postings sampled, consistent with O*NET OnLine's designation of Python as a required, growth-flagged "Hot Technology" for Statisticians, an occupation that overlaps with pharmaceutical biostatistics (www.onetonline.org). Second, generic "AI" or "machine learning" phrasing coexists with highly specific method names (RLHF, RAG, MONAI) in the same posting, which is why a keyword count of the bare term "AI" alone, the approach Indeed found dominates 74% of AI-related postings economy-wide, would undercount how technically specific pharmaceutical AI hiring has become at the senior scientist and engineer level. Third, domain-specific scientific libraries (cheminformatics, biomedical imaging) appear disproportionately in research and development (R&D) roles rather than commercial or corporate-function roles, a pattern explored further in the next section.

This sample-level detail complements, rather than replaces, Lightcast's larger indexed count for the UK Life Sciences sector, which found Artificial Intelligence mentioned in 1.5% of postings, Machine Learning in 1.2%, and smaller shares for Deep Learning (0.2%), TensorFlow (0.2%), PyTorch (0.2%), Robotics (0.1%), Generative Artificial Intelligence (0.1%), Natural Language Processing (0.1%), and the Scikit-Learn Python package (0.1%) (www.datocms-assets.com). Lightcast separately found that job postings mentioning IT-related skills such as machine learning and Python increased by roughly seven percentage points between the 2020/21 and 2023/24 periods it tracked (www.datocms-assets.com), and cited industry research from the Association of the British Pharmaceutical Industry (ABPI) finding that BioPharma roles often demand digital expertise in areas such as biomedical imaging and bioinformatics (www.datocms-assets.com), directly echoing the imaging- and cheminformatics-specific postings in Table 1.

Roles and Functions Demanding AI Skills

AI skill requirements are not evenly distributed across pharmaceutical job functions; they cluster around a recognizable set of roles that this section maps against standardized occupational data. *Table 2* organizes the roles observed in this report's posting sample and in government occupational data by function.

Table 2: Pharma Roles Carrying AI Skill Requirements, by Function

Role or Occupation	Function	Typical AI-Related Requirements	Occupational Context
AI Scientist / Applied AI Director	R&D, biomedical research	PhD in machine learning, computer science, applied mathematics, or computational biology; deep learning frameworks	Novartis's AI Scientist, Image Analysis & Digital Pathology posting requires a "Ph.D. in Machine Learning, Computer Science, Applied Mathematics, Computational Biology or related field" (novartis.wd3.myworkdayjobs.com)
AI/ML Engineer (platform and MLOps)	IT, digital platforms	8+ years AI engineering experience; cloud, CI/CD, model deployment	Novartis's remote Director, AI Foundations Engineer posting requires "8 plus years of AI engineering experience" (novartis.wd3.myworkdayjobs.com)
AI Research Scientist (generative AI focus)	R&D, technology centers	GenAI, deep learning, NLP, computer vision (CV), ML frameworks	Merck (MSD)'s AI Research Scientist posting requires "Proficiency in GenAI, deep learning, NLP, CV, and ML frameworks" (msd.wd5.myworkdayjobs.com)
Data Scientist (general and clinical)	Commercial, R&D, clinical operations	Python, SQL, statistical modeling; for clinical variants, healthcare database management	BLS classifies Data Scientists broadly, projecting 35% employment growth from 2025 to 2035 (www.bls.gov); O*NET separately profiles Clinical Data Managers as a Data Scientist specialty that applies "knowledge of health care and database management to analyze clinical data" (www.onetonline.org)
Computational/Cheminformatics Scientist	Drug discovery, pharmacokinetics	Cheminformatics libraries, molecular representations, structure-based modeling	Pfizer's Senior ML/AI Cheminformatics Scientist posting names "RDKit, OpenEye, DeepChem, InChI, SMILES/SMARTS" (pfizer.wd1.myworkdayjobs.com)
Biochemist / Biophysicist with computational skills	Discovery research	Statistical and modeling software alongside wet-lab domain knowledge	BLS reports Pharmaceutical and medicine manufacturing employs about 19% of Biochemists and Biophysicists (www.bls.gov), an occupation BLS projects to grow 6% from 2024 to 2034 (www.bls.gov)

Role or Occupation	Function	Typical AI-Related Requirements	Occupational Context
GenAI/Data Science Manager (product development)	Product development, regulatory-adjacent functions	LLMs, RAG, agentic workflow design applied to development documents and processes	Bristol Myers Squibb's Manager, Pharmaceutical Product Development GenAI & Data Science posting is scoped explicitly around " Large Language Models, Reasoning Models, Retrieval-Augmented Generation (RAG), Agentic Workflow Design " (bristolmyerssquibb.wd5.myworkdayjobs.com)

The occupational-data layer helps explain why AI skill demand concentrates where it does. The Computer and Information Research Scientists occupation, which BLS projects to grow 20% from 2024 to 2034 (www.bls.gov), draws roughly 19% of its jobs from "Research and development in the physical, engineering, and life sciences," where the median wage of \$153,430 exceeds the occupation's overall median (www.bls.gov). The median annual wage for Data Scientists was \$120,230 in May 2025 (www.bls.gov), a pattern consistent with LinkedIn's finding that AI job postings have roughly doubled since 2023, and the typical posting lists about \$177,000 in compensation, compared with \$80,000 for a non-AI role (news.linkedin.com).

McKinsey's interviews with pharma R&D leaders describe how these roles are staffed in practice: rather than hiring generalist "AI" employees, GSK "assembled a team of people who have PhD and postdoctoral qualifications in clinical imaging, alongside those with machine learning expertise" (web.archive.org), and a Johnson & Johnson leader described the target profile as "**trilingualism**" in skills proficiency, including in data science, science and medicine, and business strategy (web.archive.org). This pairing model, hybrid teams rather than hybrid individuals, is one reason a single "AI skills" job-posting keyword count understates actual AI-adjacent hiring: a domain-expert posting paired with a machine-learning-expert posting may not mention AI keywords symmetrically.

Training and Upskilling Pathways for Pharma AI Talent

A distinct question from what employers ask for is how the existing pharmaceutical workforce is meant to acquire it. Several channels have emerged, at different scales and with different degrees of formal accreditation.

Professional-society continuing education. DIA runs recurring, named courses that apply AI directly to regulated pharmaceutical functions: an "Artificial Intelligence in Pharmacovigilance" course scheduled into February 2026, and a newly introduced "**interactive workshop**" exploring how "**automation and AI are transforming signal management**" (www.diaglobal.org). The International Society for Pharmaceutical Engineering (ISPE) has built AI directly into its quality-systems curriculum: its ISPE Academy continuing-education program (ispe.org) includes a "**GAMP® Basic Principles with AI/ML**" course scheduled multiple times in late 2026 (ispe.org). The Regulatory Affairs Professionals Society (RAPS) runs a dedicated AI in Regulatory Conference, held in Barcelona in November 2026, that "**will examine how artificial intelligence is transforming regulatory affairs across the MedTech product lifecycle**" (live.raps.org), with attendance carrying 12.0 RAC or RCC recertification credits (live.raps.org) toward the credentials regulatory professionals must maintain.

University and executive-education programs. The University of California, San Francisco (UCSF) School of Pharmacy created a Master of Science in Artificial Intelligence and Computational Drug Discovery and Development (AICD3), which it describes as "**the first of its kind in the United States and the first master's**

program in the School of Pharmacy" (pharmacy.ucsf.edu), with its program director framing the goal as building "a pipeline, a bridge, between academia and industry" (pharmacy.ucsf.edu). The University of Maryland's School of Pharmacy launched a parallel Master's in AI for Drug Development in 2025, citing AI's **"transformative"** role in enhancing pharmaceutical research capabilities (www.umaryland.edu). At the executive-education level, MIT offers, through edX, a six-week course on **"the interconnectedness between the science of drug discovery, machine learning, and business decision making"**, with a session starting October 21, 2026 (www.edx.org), and Novartis co-developed a public Coursera specialization on drug discovery that had 12,890 learners already enrolled at the time of access (www.coursera.org).

Internal corporate academies. Beyond the widely reported AI-literacy programs at other large drugmakers, Novartis operates a named internal function, "Ready for AI" (RAI), whose lead role is explicitly responsible for **"the US Commercial Responsible Data and AI learning and activation program, ensuring associates and leaders understand responsible data and AI principles"** (www.novartis.com). AstraZeneca's "Thriving in the Age of AI" program had **"over 12,000 employees"** complete tiered Bronze, Silver, and Gold-level AI certifications as of April 2025 (www.astrazeneca.com), a figure the company reported had grown to **over 37,700 employees** participating by October 2025, in more than ten languages (www.astrazeneca.com), with the company noting it is **"evolving even further this year with fresh content, new experiences and advanced levels of accreditation"** in response to the shift toward agentic AI (www.astrazeneca.com).

Whether this training investment translates into measurable productivity is a separate, still-thin evidence base. EY's global Work Reimagined research found that employees receiving over 81 hours of **annual AI training** reported an average productivity gain of 14 hours per week, well above employees with less training (www.ey.com), but the same survey found that these more heavily trained employees are also 55% more likely to leave their organization (www.ey.com), a retention risk directly relevant to employers making large per-employee training investments. EY estimated that only 12% of employees economy-wide currently receive AI training sufficient to unlock the technology's full productivity benefit (www.ey.com), a figure that is not pharma-specific but frames the scale of the gap corporate academies are attempting to close.

Data Analysis and Evidence

Quantifying pharmaceutical AI-skill demand precisely requires triangulating several data sources, since no single tracker publishes a comprehensive pharma-only AI-skills index. *Table 3* collects the key measured figures cited throughout this report, alongside their scope and observation date, so the boundaries of each figure are explicit.

Table 3: Quantitative Benchmarks on AI Skills in Job Postings, by Source and Scope

Metric	Value	Scope	Source and As-Of Date
Share of postings mentioning AI-related skills	Just over 2% (Life Sciences) vs. less than 1% (all sectors)	UK Life Sciences vs. UK labor market, April 2023 to March 2024	Lightcast (www.datocms-assets.com)
Share of postings mentioning AI skills, health sector	Less than 1% (lowest of tracked sectors)	Global, health sector	PwC 2026 AI Jobs Barometer (www.pwc.com)
Growth of AI-skill-specific job postings vs. total jobs market	69% vs. 9% (about 8x faster)	Global, all sectors	PwC 2026 AI Jobs Barometer (www.pwc.com)

Metric	Value	Scope	Source and As-Of Date
"AI-touched" job titles	822 titles (Q1 2026) vs. 264 titles (Q1 2022)	US labor market, all sectors	Indeed Hiring Lab (hiringlab.indeed.com)
Barrier to AI adoption citing staff skills	55% of respondents (top-cited barrier)	Global pharma/biotech survey, 2019	Tufts CSDD / DIA (www.ebi.ac.uk)
Planned AI-related staffing increases	Nearly 60% of respondents, within 1 to 2 years	Global pharma/biotech survey, 2019	Tufts CSDD / DIA (www.ebi.ac.uk)
Life sciences leaders who scaled AI successfully	22% (only 9% report significant returns)	Global life sciences executives, survey fielded Sept to Nov 2025	Deloitte 2026 Life Sciences Outlook (www.deloitte.com)
Biopharma R&D executives reporting increased lab throughput	53% of 104 surveyed executives	US biopharma R&D executives, 2025	Deloitte Center for Health Solutions (www.deloitte.com)
Biopharma manufacturing employment growth, 2015 to 2024	24% (about 6x manufacturing overall)	US biopharmaceutical manufacturing	PhRMA / TEconomy (phrma.org)
Projected annual biopharma hiring/reskilling need through 2034	Roughly 30,000 workers per year	US biopharmaceutical industry	PhRMA / TEconomy (phrma.org)
Data Scientist employment growth, 2025 to 2035	35%	US, all industries	BLS Occupational Outlook Handbook (www.bls.gov)

Two patterns emerge from placing these figures side by side. First, posting-level keyword prevalence (Lightcast's 2%, PwC's under-1% health figure) and workforce-planning sentiment (Tufts CSDD/DIA's 55% citing staff skills as a barrier, PhRMA's 30,000-worker annual reskilling estimate) point in the same direction from different angles: AI-specific keywords are still rare in the typical posting, but the roles and skills gaps behind that rarity are already large enough to worry survey respondents and to justify economy-scale training investment. Second, the growth-rate figures are consistently steep even where the base is small. PwC's finding that AI-skill-specific postings are growing roughly eight times faster than the overall jobs market (www.pwc.com) is echoed by Indeed's more than three-fold rise in AI-touched job titles between 2022 and 2026 (hiringlab.indeed.com). Separate labor-market research from Revelio Labs finds that firms identified as AI adopters, via AI-integrator job postings, grew headcount 27% more than non-adopters since October 2022, with the gain concentrated in senior roles (www.reveliolabs.com), a pattern consistent with pharma's own hybrid-team hiring model described above. None of these growth-rate figures should be read as pharma-specific unless explicitly labeled as such in Table 3; where a figure is economy-wide, this report presents it as context for, not a measurement of, pharmaceutical hiring specifically.

Deloitte's 2026 outlook adds a workforce-productivity angle directly relevant to training design: 29% of biopharma leaders and 31% of medtech leaders said they plan to use AI tools or training specifically to improve workforce productivity in 2026 (www.deloitte.com), while only 14% report full AI-workflow implementation today, with another 40% still working toward it (www.deloitte.com). A companion IntuitionLabs analysis of pharma's broader AI skills gap covers the underlying reskilling-versus-hiring economics and named corporate literacy-program scale in more detail (intuitionlabs.ai); the figures assembled here are narrower, focused specifically on posting-level skill language and role taxonomy rather than survey sentiment about the gap itself.

Case Studies and Real-World Examples

Three neutral, independently documented examples illustrate how pharmaceutical organizations are operationalizing AI-skill acquisition, each addressing a different layer of the talent pipeline described above.

AstraZeneca: Enterprise-Wide Tiered AI Accreditation

AstraZeneca's "Thriving in the Age of AI" initiative applies a graduated Bronze, Silver, Gold (and later Platinum, Diamond) accreditation structure across its global workforce, rather than a single one-time training event. The company reported over 12,000 employees had achieved at least one certification tier by April 2025 (www.astrazeneca.com), a figure it reported had grown to over 37,700 participants by October 2025 across more than ten languages (www.astrazeneca.com). The company has framed the program as continuing to expand rather than concluding after its initial rollout, stating it is **"evolving even further this year with fresh content, new experiences and advanced levels of accreditation"** in light of the shift toward agentic AI (www.astrazeneca.com), a case of training accreditation being treated as an ongoing program rather than a one-time compliance exercise.

UCSF and University of Maryland: A New Academic Track for Computational Drug Development

Where corporate academies retrain existing staff, two universities have built new degree tracks aimed at producing AI-literate pharmaceutical scientists from the outset. UCSF's School of Pharmacy created its Master of Science in AI and Computational Drug Discovery and Development in 2024, describing it as the first program of its kind in the United States (pharmacy.ucsf.edu). The University of Maryland followed in 2025 with its own Master's in AI for Drug Development (www.umaryland.edu). Both programs sit alongside, rather than replace, executive-education offerings such as MIT's AI in Pharma and Biotech course delivered through edX (www.edx.org) and Novartis's public Coursera specialization (www.coursera.org), suggesting the talent pipeline is being built at multiple altitudes simultaneously: degree programs for new entrants, executive courses for mid-career managers, and free or low-cost specializations for a broader public audience.

Novartis: A Named Internal Function for Responsible AI Learning

Rather than treating AI training as an ad hoc project, Novartis created a standing internal role and program, "Ready for AI" (RAI), with a mandate covering **"the US Commercial Responsible Data and AI learning and activation program"** (www.novartis.com). Structuring AI literacy as a permanent program with a named accountable role, rather than a one-time training rollout, is a governance pattern that a life-sciences AI advisory would recognize: IntuitionLabs describes its own implementation approach in similar terms, prioritizing **"governed workflows"** and measuring **"adoption, time recovered, quality, and support before scaling"** (intuitionlabs.ai) rather than a single enterprise-wide rollout.

Implications and Future Directions

Several implications follow from combining posting-level, occupational, and training-pipeline evidence. First, the low headline percentages (Lightcast's just-over-2% for UK Life Sciences, PwC's under-1% for the health sector) should not be read as evidence that AI skills are marginal in pharma; the direct posting sample in Table 1 shows that where AI skills do appear, they are highly specific, spanning generative AI, reinforcement learning, and domain-specific scientific tooling, not superficial keyword insertion. Employers appear to be concentrating specific, hard-to-fill AI skill requirements into a relatively small number of senior scientist and engineering postings rather than diffusing generic "AI familiarity" language across the broader posting base, consistent with Indeed's finding that non-tech titles now account for the majority of AI-touched titles economy-wide even as the specific technical language within pharma R&D postings remains dense (hiringlab.indeed.com).

Second, the training ecosystem is bifurcating by design rather than by accident. Enterprise-wide accreditation programs, such as AstraZeneca's, aim at broad AI literacy across a workforce that will mostly use AI tools rather than build them, while professional-society courses (DIA's pharmacovigilance AI course, ISPE's GAMP AI/ML training, RAPS's regulatory AI conference) target the much narrower population that must certify AI use against regulatory and quality frameworks. New university degree programs target a third population entirely: graduates entering the field already fluent in both pharmaceutical science and computational methods. A hiring manager evaluating candidates for the roles in Table 2 should expect to draw from all three pipelines simultaneously, since none alone produces the "trilingual" profile McKinsey's sources described (web.archive.org).

Third, gender composition in AI hiring is a talent-supply constraint worth tracking alongside skills taxonomy. LinkedIn's research found women accounted for just 26% of AI hires in 2025, against 50% of hires into non-AI occupations (news.linkedin.com), a broader labor-market pattern that, if it holds within pharma specifically, would narrow the effective applicant pool for the roles in Table 2 beyond what raw postings volume suggests. LinkedIn's own "Jobs on the Rise" research also found that 56% of professionals plan to job-hunt in 2026, while 76% say they do not feel prepared for an AI-reshaped job market (www.linkedin.com), a readiness gap that pharma employers competing for the same cross-industry AI talent pool will need to account for in how they write and target postings.

Looking ahead, the clearest reproducible extension of this report's method would be a larger, systematically sampled posting dataset, drawn consistently across a defined set of drugmakers on a fixed cadence, classified against the same schema used in Table 1. That would convert the illustrative sample presented here into a genuine time series capable of tracking whether pharma-specific AI-skill prevalence is converging toward, or diverging from, the broader life-sciences and economy-wide benchmarks in Table 3.

Frequently Asked Questions (FAQs)

What counts as an "AI skill" in a pharma job posting? This report classifies a posting as containing an AI skill only when it names a specific machine-learning method, model family, framework, or tool (Python paired with a deep-learning framework, LLMs, RAG, reinforcement learning, or a domain-specific library such as RDKit), rather than the bare word "AI," since Indeed's Hiring Lab found nearly 74% of AI-related postings used only the generic term (hiringlab.indeed.com).

Which pharma job postings are most likely to require machine learning skills specifically? In the postings sampled for this report, machine learning and deep-learning framework requirements (Python, PyTorch, TensorFlow) concentrated in R&D-adjacent titles such as AI Scientist, Data Scientist (biomedical imaging or discovery-focused), and AI/ML Engineering roles, rather than in commercial or general corporate-function postings, consistent with BLS data showing a median annual wage of \$120,230 for Data Scientists in May 2025 (www.bls.gov).

What data science skills are pharmaceutical companies hiring for beyond generic "AI"? Beyond Python and SQL, current postings named cheminformatics libraries (RDKit, OpenEye, DeepChem), biomedical imaging tools (MONAI, Hugging Face, Lightning), generative AI components (LLMs, RAG, agentic workflow design), and MLOps/cloud tooling (AWS, MLFlow, Docker, CI/CD), as detailed in Table 1.

Are AI skills required for regulatory affairs or pharmacovigilance roles? Not typically as a baseline hiring requirement, but professional continuing-education bodies are actively building AI into these functions' credentialing pathways: DIA runs recurring courses applying AI and NLP to pharmacovigilance signal management (www.diaglobal.org), and RAPS offers a dedicated AI in Regulatory Conference carrying recertification credit (live.raps.org), signaling that AI literacy is becoming an expected competency for these credentials even where it is not yet a hiring-line requirement.

How is the pharmaceutical industry addressing its AI skills gap? Through three parallel channels documented in this report: enterprise-wide internal accreditation programs (AstraZeneca's tiered certification, Novartis's Ready for AI program), new university degree tracks (UCSF's and the University of Maryland's AI-focused pharmacy master's programs), and public or semi-public online specializations (Novartis's Coursera program, MIT's edX executive course). A companion IntuitionLabs report examines the reskilling-versus-hiring economics behind these investments in more depth (intuitionlabs.ai).

What AI hiring trends are visible across life sciences companies more broadly? Growth-rate indicators point in one direction even where base rates remain low: PwC found AI-skill-specific job postings growing roughly eight times faster than the overall jobs market (www.pwc.com), and Revelio Labs found AI-adopting firms grew headcount 27% faster than non-adopters (www.reveliolabs.com). For company-by-company hiring, layoff, and compensation signals at large drugmakers, a companion IntuitionLabs census covers that ground directly (intuitionlabs.ai).

Conclusion

AI-related keywords remain a minority requirement in pharmaceutical job postings by every large-sample measure available, from Lightcast's just-over-2% share of UK Life Sciences postings to PwC's under-1% figure for the health sector globally. Read alongside a direct sample of current AI and machine-learning postings from six major drugmakers, however, that low prevalence coexists with unusually specific technical demands wherever AI skills do appear: named frameworks, generative-AI components, reinforcement learning techniques, and domain-specific scientific libraries, concentrated in R&D, cheminformatics, and platform-engineering roles rather than spread evenly across the organization.

Government occupational data situates these roles within faster-than-average growth projections, and a widening set of professional-society courses, university degree programs, and enterprise accreditation initiatives has emerged to supply them. None of these channels alone appears sufficient: survey evidence from Tufts CSDD and DIA as far back as 2019 already identified staff skills as the leading barrier to AI adoption in biopharmaceutical development, and PhRMA's more recent workforce projections put the industry's annual hiring and reskilling need at roughly 30,000 workers through 2034. The practical implication for hiring managers, career-changers, and workforce planners alike is the same: treat "AI skills" as a defined, classifiable set of named tools and methods tied to specific roles, not a single generic requirement, and expect to draw simultaneously from internal upskilling, professional certification, and new academic pipelines to fill it.

For informational purposes. Verify critical medical, legal, financial, regulatory, and technical information with qualified professionals and primary sources.