

Data Scientist

Location: Remote - Dublin, Ireland

Reporting to: Head of Product (US-based)

Team: Product

Supply Wisdom is a global leader in transformative risk intelligence, offering real-time insights to drive business growth, reduce costs, enhance security and compliance, and identify revenue opportunities. Our AI-based SaaS products cover various risk domains including financial, cyber, operational, ESG, and compliance. With a diverse workforce that is 57% female, our clients include Fortune 100 and Global 2000 firms in sectors like financial services, insurance, healthcare, and technology.

The Role

Supply Wisdom is hiring a Data Scientist to join our Product team, working directly on the models, pipelines, and data infrastructure that power our risk intelligence platform. You'll join our existing data science function and report to the Head of Product, contributing across everything from predictive risk models to production data pipelines to customer-facing analytical questions.

This is a hands-on, individual-contributor role for someone who can take a problem, own it end to end, and drive it to a working outcome without needing to be managed step by step. You'll move fluidly between statistical modeling, ML engineering, and data pipeline work, exercising judgment about which approach best fits the problem rather than defaulting to one tool. We're looking for someone who treats delivery as iterative rather than a single big push to get it perfect.

You'll be joining a growing local team in Ireland. You may occasionally be asked to join in-person work sessions at a local workspace.

What You'll Do

- Take on ambiguous, open-ended problems (e.g. "improve target coverage in this risk domain" or "reduce false positive rate for this classifier") and independently structure an approach, build it, and iterate toward a working solution.
- Design, build, and maintain data pipelines and ML models that identify, detect, and quantify risk intelligence across financial, cyber, operational, ESG, and compliance domains.
- Prepare, clean, and structure large and often messy datasets for modeling, exercising judgment on where automation, direct data integration, or LLM-based approaches each make the most sense.
- Build and continuously refine predictive and classification models (e.g. credibility scoring, urgency/severity classification, entity resolution), evaluating performance against real outcomes and iterating based on data-driven feedback.

- Engage directly with product, engineering, and occasionally customer-facing stakeholders to translate business and methodology questions (e.g. model bias, data confidence, coverage limitations) into clear technical answers and solutions.
- Use Python and standard data science/ML libraries alongside strong database and querying skills to move fluidly from data prep to modeling to production.
- Proactively evaluate and adopt new tools and techniques, including AI-assisted workflows, to accelerate your own delivery rather than defaulting to manual or established methods.
- Document your methodology, code, and data schemas clearly enough that teammates and stakeholders can build on your work.

What Sets You Apart

Beyond the technical skill set, this role calls for a specific working style:

- Self-driven: given a goal, you figure out the path there. You don't need the problem broken into sub-tasks before you can start.
- Outcome-oriented and iterative: you'd rather ship a working first version, get it in front of real data or feedback, and revise than spend extra time trying to perfect the approach before anything ships.
- Comfortable with ambiguity: you can take a rough problem statement, ask the right clarifying questions, and still move forward while some details are unresolved.
- Genuinely curious about new tools and methods, and willing to experiment with them on your own initiative rather than waiting to be told to try something new.

What You Bring

- 3-5 years of experience in applied data science, machine learning, or a closely related technical role.
- Strong Python skills across the data science and ML stack (pandas, NumPy, scikit-learn, TensorFlow, PyTorch or equivalent).
- Solid database and data engineering fundamentals — comfortable designing schemas, writing efficient queries, and building reliable pipelines, not just consuming clean data.
- Experience deploying models into production and monitoring their performance, not just building them in a notebook.
- Excellent written and verbal communication skills; able to explain technical tradeoffs to non-technical stakeholders.
- A track record of working with minimal supervision and following through on open items without prompting.
- Degree in Computer Science, Applied Mathematics, Statistics, or a related field, or equivalent practical experience.
- Nice to have: Exposure to REST API design and development (Django Rest Framework or similar).

Our Commitment to You

We offer a competitive salary and generous benefits, along with a vibrant work environment and a global team filled with passionate people from diverse cultures and backgrounds.

If you're looking to make an impact delivering market-leading risk management solutions, empowering our clients, and doing work that matters, Supply Wisdom is the place for you.

Learn more at supplywisdom.com and on [LinkedIn](#).