JUPITER PRIME

Consulting • Training • Advisory

جوبيتر برايم للاستشارات والتدريب

COURSE BROCHURE

Human-Centred Machine Learning Product Design Course

3 – 7 May 2027
Vienna

jupiterprime.uk

info@jupiterprime.uk

Registered in the United Kingdom and Saudi Arabia

Human-Centred Machine Learning Product Design Course

Ref: 1435_53473 **Date:** 3 – 7 May 2027 **Location:** Vienna **Fees:** 6600 GBP

Human-Centred Machine Learning Product Design is a 1-week training course for product managers, UX leads and data team leads who decide how a prediction or recommendation reaches its users, and it produces a User-Centred Model Specification for one feature. Models now shape loan offers, service queues and clinical alerts, yet many fail because nobody asked the people who act on their output. Jupiter Prime Training runs this course on Human-Centred Machine Learning Product Design.

Why Human-Centred Machine Learning Product Design Matters Now

Three forces push design to the front of machine learning work. Users reject predictions they cannot question, so accuracy alone no longer secures uptake. Domain experts hold the judgement that good labels and sensible thresholds depend on, yet are often consulted last. People affected by automated decisions expect fair treatment and a way to contest outcomes. Products that involve users from framing to feedback earn trust and steady use; products built only around a metric are bypassed, overridden or quietly abandoned. The course uses the NIST AI Risk Management Framework AI RMF and ISO/IEC 42001 as reference points for human-centred design duties.

Who Gains Most from Human-Centred Machine Learning Product Design

- Product leads who decide which user problem a model should address and how success is judged
- UX researchers and designers who shape how predictions, confidence and explanations appear on screen
- Data team leads who agree labelling rules and thresholds with domain experts
- Service owners who handle complaints and appeals about automated outcomes
- Analysts who run usability and fairness evaluations before release

A product squad gains most when its product lead, designer and data lead attend together and share one specification.

Human-Centred Machine Learning Product Design Capabilities

Participants leave able to put users and domain experts at each point where a model's design choices are made.

- Design user research that frames a prediction problem around real decisions — product roadmap choices
- Specify labelling guidelines with domain experts to reduce ambiguous examples — data quality owners
- Build explanation and confidence displays matched to user expertise — frontline decision makers
- Test model outputs with affected user groups for unequal error impact — fairness and ethics reviewers
- Apply usability evaluation methods to prediction interfaces — release approval
- Design feedback loops that capture corrections and appeals — model retraining priorities

Key Human-Centred Machine Learning Product Design Terms

- Problem framing — defining the user decision a model supports before choosing data, target or metric.
- Labelling guideline — written rules that tell annotators how to classify difficult or borderline examples consistently.
- Mental model — the user's own picture of how a system works, which shapes how they read its outputs.
- Explanation interface — the on-screen element that shows why a model produced a given prediction.
- Contestability — the ability of an affected person to question, correct or appeal an automated outcome.
- Feedback loop — a route by which user corrections return to improve data, labels or model behaviour.

Human-Centred Machine Learning Product Design Day by Day

Day 1: Users, Decisions and Model Purpose

- Contextual Inquiry Interviews — Observing Decisions a Model Might Support
- Problem Framing Canvas — Choosing Target, Metric and Intended Use
- Stakeholder and Affected Group Mapping — Who Gains or Bears Risk
- NIST AI RMF Map Function — Recording Context and Intended Users
- Retail Recommendation Case — When Accuracy Did Not Bring Uptake

Day 2: Data, Labels and Domain Expertise

- Labelling Guideline Design — Rules for Borderline and Ambiguous Examples
- Expert Annotation Sessions — Capturing Tacit Judgement From Specialists
- Inter-Annotator Agreement Checks — Spotting Unclear Categories Early
- Representative Sampling Plans — Including Under-Served User Groups
- Threshold Setting Sessions With Experts — Balancing Misses and False Alarms

Day 3: Explanations and Interaction Design

- Mental Model Elicitation — How Users Believe the System Works
- Explanation Interface Patterns — Reasons, Examples and Counterfactuals
- Confidence Display Options — Scores, Bands and Uncertainty Cues
- Override and Escalation Controls — Keeping Users in Charge
- Clinical Alert Case — Designing Explanations for Busy Nurses

Day 4: Fairness, Usability and Feedback

- Fairness Testing With Affected Users — Unequal Error Experiences
- Usability Evaluation of Model Outputs — Task Success and Trust Ratings
- Contestability and Appeal Routes — Handling Disputed Automated Outcomes
- Feedback Loop Design — Turning Corrections Into Training Signals
- ISO/IEC 42001 Impact Assessment Inputs — Documenting Human-Centred Evidence

Day 5: Practising the User-Centred Model Specification

- Research Synthesis Session — Applying Contextual Inquiry Findings to One Feature
- Label and Threshold Agreement — Applying Expert Annotation Results
- Interface Prototype Critique — Applying Explanation and Confidence Patterns
- Evaluation Plan Drafting — Fairness and Usability Tests Before Release
- User-Centred Model Specification — Completing Purpose, Labels, Interface and Feedback

Applying Human-Centred Machine Learning Product Design After the Course

- Run five contextual inquiry interviews with users of an existing model and log where its outputs are ignored.
- Hold one expert annotation session to rewrite the labelling guideline for the most disputed category.

Human-Centred Machine Learning Product Design Questions FAQ

What is Human-Centred Machine Learning Product Design and why does it matter now?

It is the practice of shaping a model's purpose, labels, interface and feedback around the people who use or are affected by it. It matters now because users ignore or override predictions they cannot question.

Which roles gain most from Human-Centred Machine Learning Product Design?

Product leads, UX researchers and designers, data team leads, service owners and evaluation analysts gain most.

How does Human-Centred Machine Learning Product Design differ from a technical machine learning course?

It focuses on users rather than algorithms. A technical course covers model building and tuning; this course covers framing, expert labelling, explanation interfaces and testing with affected people.

What do you bring back from Human-Centred Machine Learning Product Design?

You bring back a User-Centred Model Specification for one real or planned machine learning feature.

Human-Centred Machine Learning Product Design: The Way Forward

The User-Centred Model Specification gives product, design and data leads one shared statement of who a model serves, how it is labelled, how it explains itself and how users correct it. Its first use is usually a go or no-go decision on a feature in development. Revisit the explanation design after release, when real usage shows which users still override the model.