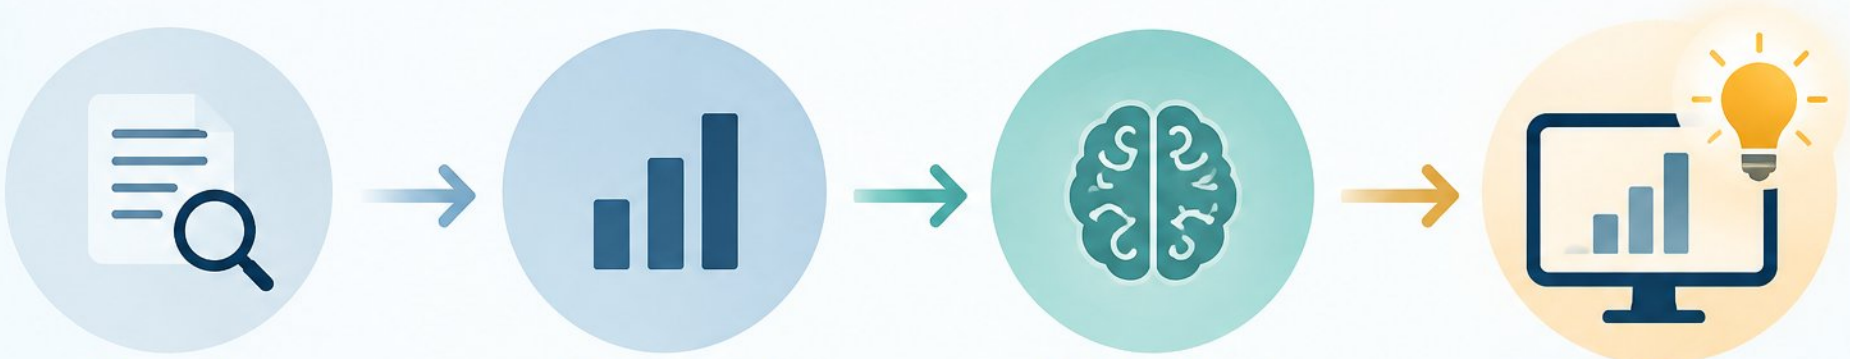
# Machine Learning in Business Intelligence

- How ML fits inside BI workflows and data products
- For BI analysts, analytics engineers, data-product owners, and managers
- Spot high-value ML opportunities and build data-ready BI pipelines
- Evaluate, explain, and govern ML-in-BI initiatives
- Why now: embedded, predictive, decision-support experiences go mainstream
- Concepts, use cases, data readiness, architecture, delivery, trust, governance, 90-day plan



# From Dashboards to Embedded Intelligence




Descriptive, predictive, prescriptive: three levels of dashboard intelligence

ML wins on forecasting, scoring, ranking, and anomaly detection

Workflows shift: new handoffs, feature work, governed model outputs, monitoring duties

Blockers: inconsistent metrics, unclear ownership, skills gaps between BI and data science

# Core *ML* Concepts in BI Language

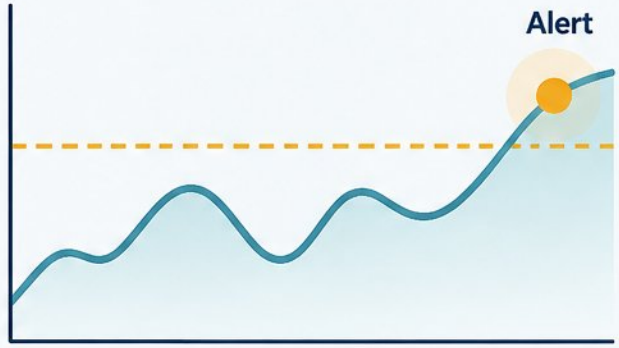
DATA
INSIGHTS
BETTER DECISIONS

- Features and labels map to familiar BI entities: orders, accounts, periods
- Supervised predicts a known target; unsupervised segments; time-series forecasts
- Metrics tie to decisions: precision, recall, MAE, lift, calibration, thresholds
- Watch for overfitting, drift, and leakage in warehouse-backed pipelines
- Always clear a rule-of-thumb baseline before scaling









# High-Value Use Cases by Function

## Finance

Forecasting



Forecasting: regression models publish demand, revenue, and pipeline projections to planning dashboards.

## Sales

Churn scoring



Churn scoring: classification flags at-risk customers; value lies in the retention action it triggers.

## Operations

Anomaly detection



Anomaly detection: surface operational, financial, and pipeline issues where analysts already review KPIs.

## Marketing

Segmentation, propensity, and next-best-action



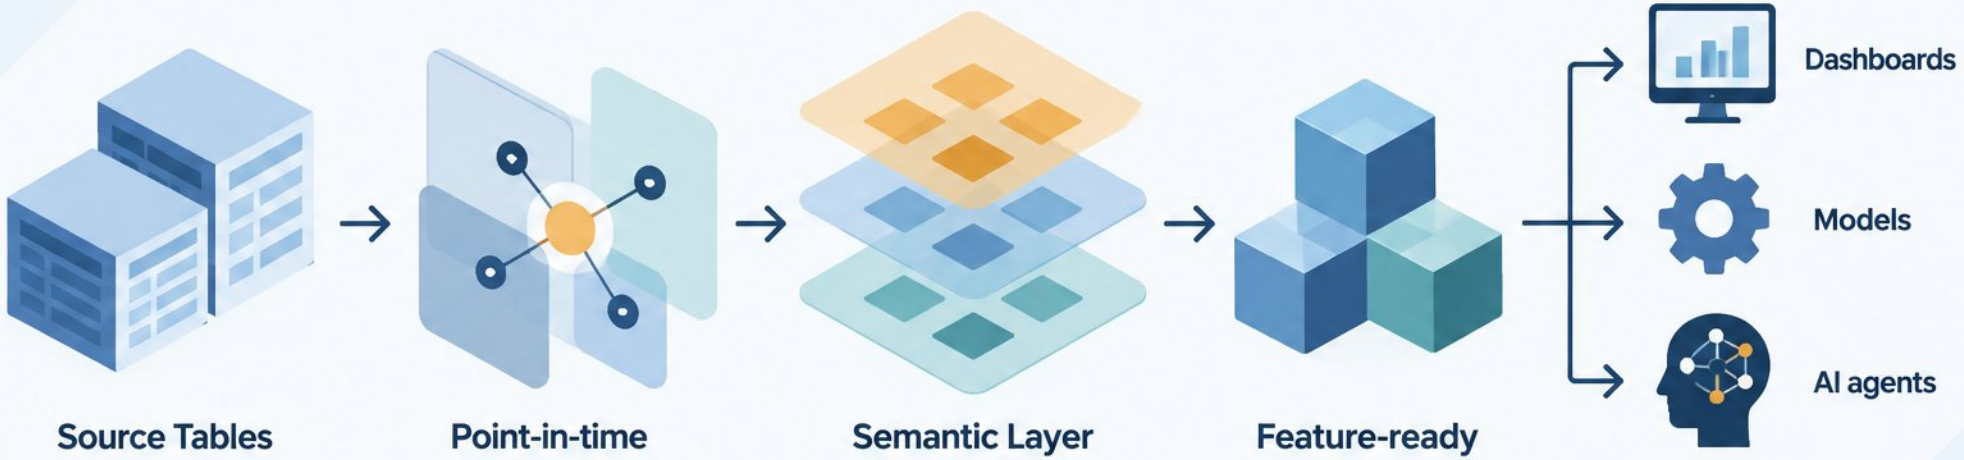
Segmentation, propensity, and next-best-action embedded in BI surfaces and workflows.

**Prioritize by business value, feasibility, data readiness, and time to measurable outcome.**

# Data Readiness and the Semantic Layer

From trusted data to broader impact

CLEAN DATA.
SHARED MEANING.
BIGGER POSSIBILITIES.




Completeness, freshness, grain, lineage, and stable definitions over time

Semantic layer defines metrics once, reused by dashboards, models, and AI agents

Feature-ready marts with point-in-time correct joins prevent data leakage

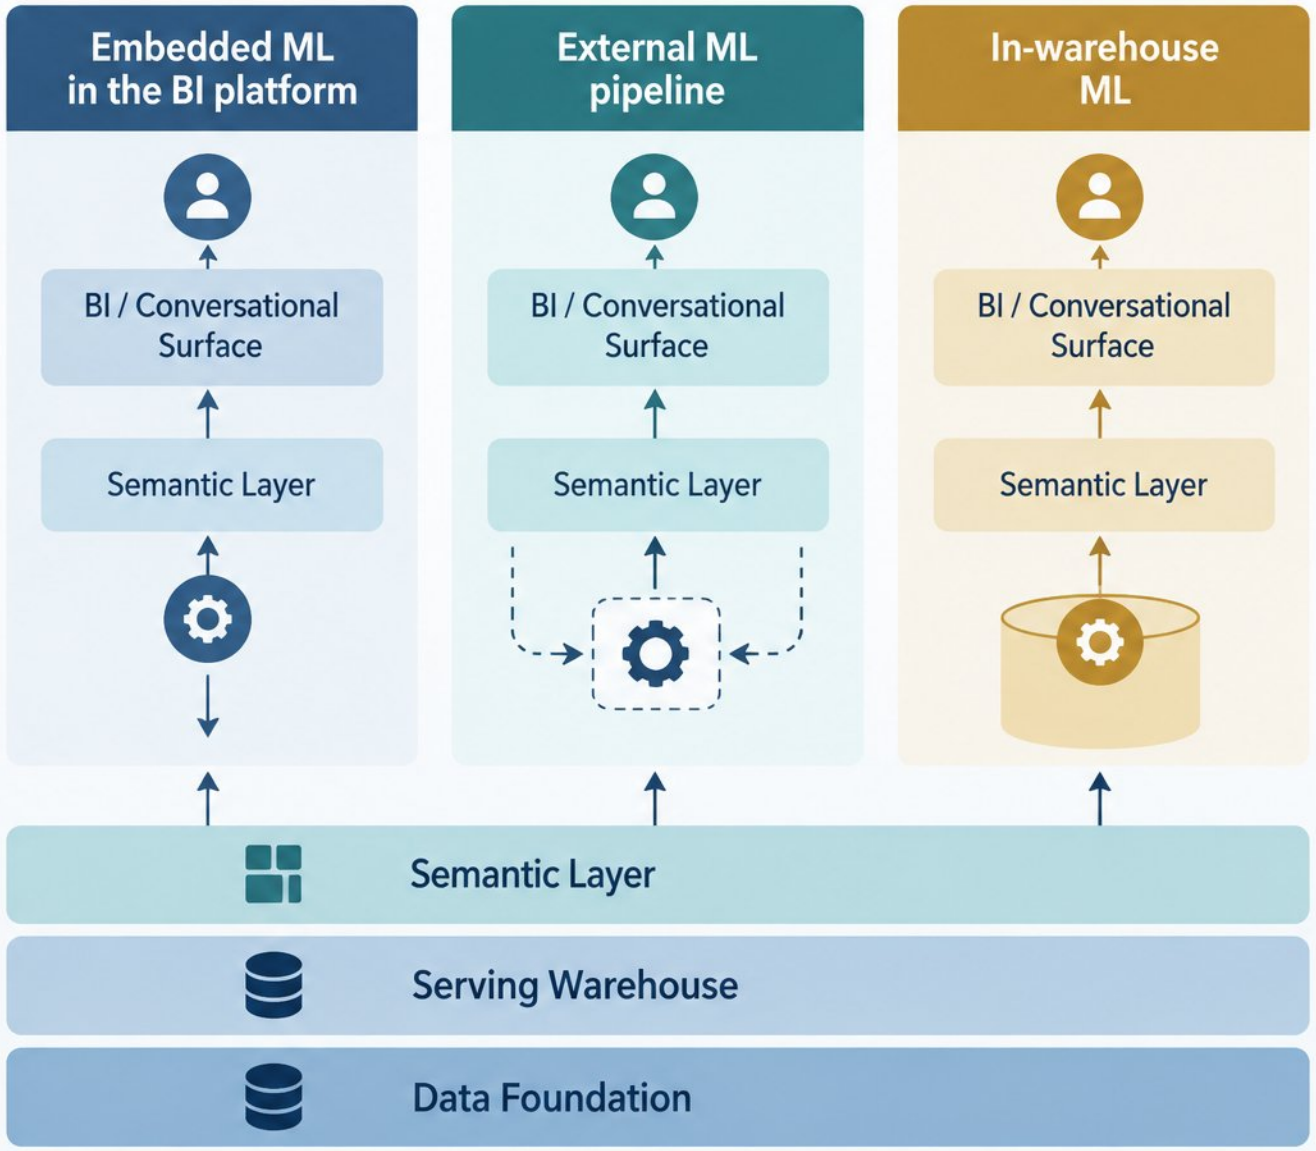
Clear ownership: data engineers, analytics engineers, and BI analysts

Minimum data-quality checks before handing data to any model

# Tooling Landscape and Architecture Patterns

PEOPLE / DATA / MODELS / INSIGHTS

- Four-layer stack: data foundation, serving warehouse, semantic layer, BI/conversational surface
- Embedded ML in BI vs external pipelines vs in-warehouse ML — each has trade-offs
- Warehouse-native keeps data fresh; caching and pre-aggregation control cost and latency
- Choose batch, real-time, or in-dashboard scoring by freshness, cost, and complexity
- Build vs buy vs platform-native: match choice to team skills and time to value



# Delivery Workflow: Question to Model to Dashboard

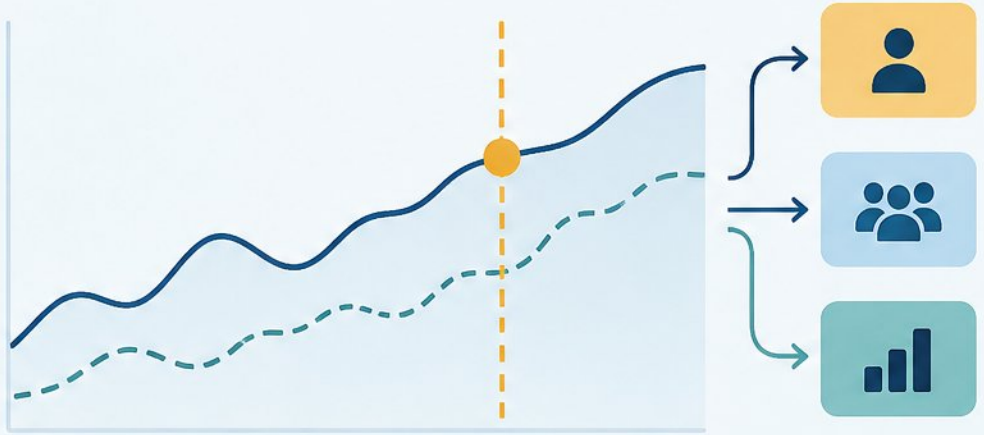


| Business | Data Science | Analytics Engineering | BI | Monitoring |
| --- | --- | --- | --- | --- |
| Frame the business question as an ML problem with a clear decision it changes | Validate against simple rules and a rule-of-thumb baseline before advanced modeling | Assign ownership: data science for models, analytics engineering for features, BI for surface | Ship model outputs as governed metrics and data products in existing BI experiences | Monitor performance, define retraining triggers, and close the feedback loop |



# Evaluation and Business Alignment

- Align trade-offs to business KPIs: precision vs coverage, recall vs outreach cost
- Measure lift against a baseline, framed in revenue or retention terms
- Translate precision, recall, MAE, and thresholds into the downstream action
- Favor reliability: an 85% model that is always available beats a 92% model that fails
- Evaluate by segment to reveal where the model quietly underperforms





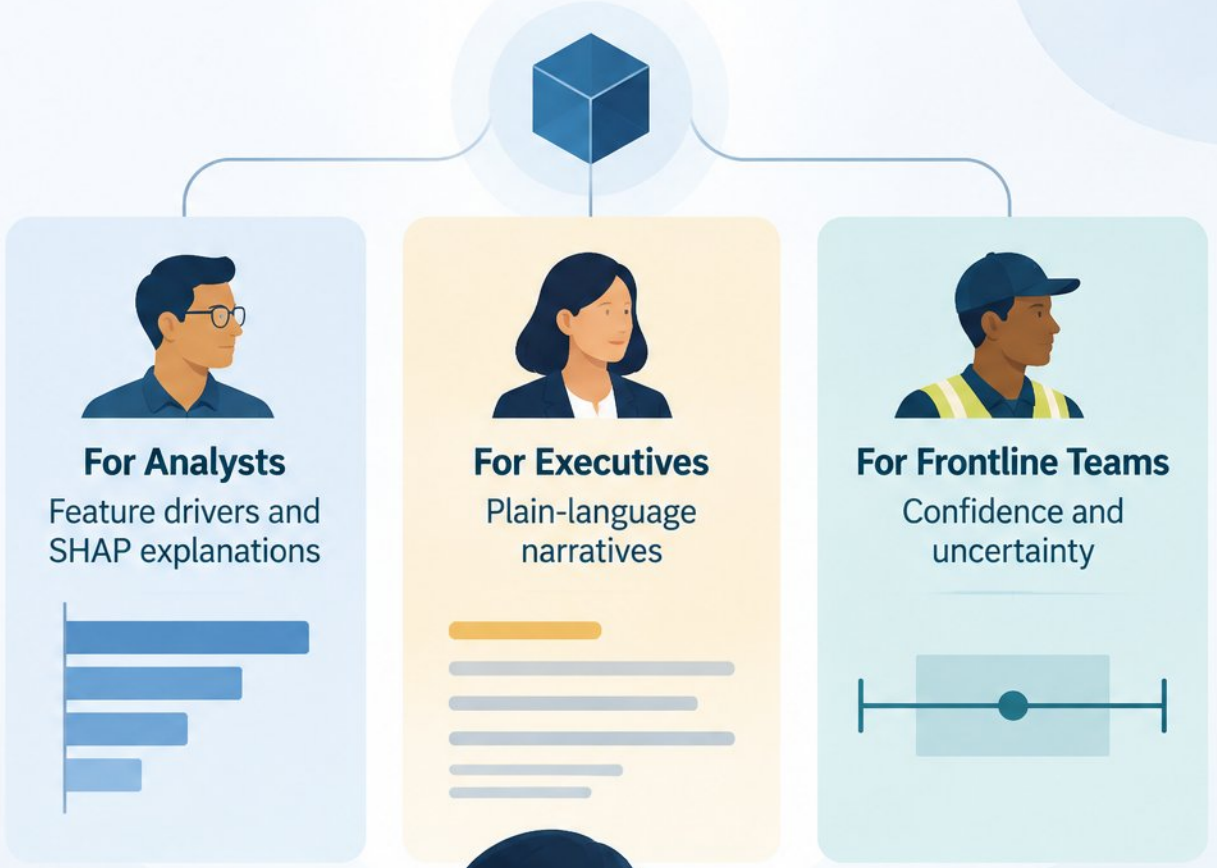
Precision vs Coverage



Recall vs Outreach Cost

# Explainability and Communicating Uncertainty

- Tailor explanations to executives, analysts, and frontline teams by technical depth.
- Use feature drivers, SHAP score explanations, and plain-language narratives in BI surfaces.
- Show confidence and uncertainty without implying misleading precision.
- Surface common failure modes visibly in dashboards rather than hiding them.
- Design contestability and human oversight in from the start, not as afterthoughts.






# Governance, Risk, and Responsible ML

TRUST
RESPONSIBLY
BUILD TOMORROW

- Bias and fairness reviews when scores drive customer or operational decisions
- Privacy, consent, and row-level access controls aligned with the warehouse
- Model documentation, versioning, rollback, and audit trails for every deployed score
- Regulatory anchors: EU AI Act, NIST AI RMF, ISO/IEC 42001, FSB sound practices
- AI system and model registry, named human oversight, and periodic review



# Team Roles, Skills, and RACI

PEOPLE / SKILLS / ALIGNMENT FOR IMPACT

| Initiative / Activity | BI analysts | analytics engineers | data scientists | product owners | managers |
|---|---|---|---|---|---|
| Frame the decision | C | C | C | R | A |
| Design solution | R | A | C | C | I |
| Build and test | C | R | A | I | I |
| Drive adoption | C | C | C | R | A |
| Monitor and iterate | R | C | A | C | I |

**Analyst skill roadmap**




- RACI spans BI analysts, analytics engineers, data scientists, product owners, managers
- Analyst skill roadmap: statistics literacy, feature design, evaluation, model monitoring
- Product owners frame the decision, own adoption, define success metrics
- Managers sequence funding, accept risk, remove organizational blockers
- Pair new skills with real use cases so learning sticks

# Measuring Adoption and Business Impact

From model output to measurable value

## Adoption metrics

- Adoption metrics: dashboard usage, score consumption, action rate, time-to-decision

## Impact metrics

- Impact metrics: revenue retained, forecast reliability, downtime reduced, working capital




- Review cadence: weekly model oversight, periodic performance and system reviews

- Report ML-in-BI value upward through the BI surfaces teams already trust

# 90-Day Action Plan and Playbook

FROM IDEA TO IMPACT, WITH A HUMAN IN THE LOOP

PEOPLE
INSIGHTS
ACTIONS
BETTER OUTCOMES

### Weeks 1–3

pick one use case, confirm the decision, define success

### Weeks 4–7

prepare data via the semantic layer, build a rule-based baseline

### Weeks 8–11

ship a pilot score with SHAP-style explanations and an intervention playbook

### Weeks 12–13

review adoption and impact, document the model, decide scale, iterate, or stop

**Avoid:** skipping the baseline, ignoring drift, shipping scores without an action owner



# Key Takeaways and Discussion



- Winning pattern: one high-value use case, governed semantic layer, measurable outcome in 3–6 months
- Reliability, trust, and governance differentiate success, not model sophistication
- BI teams own the most valuable asset: trusted metrics and business context
- Next steps: pick a pilot, assign a decision owner, schedule baseline review
- Discussion and self-assessment: what will you ship in 90 days?

PersonWise.

PersonWise • AI digital-human course platform



Scan the code or click anywhere to watch this course live, with voice Q&A

personwise.ai/t/b72b73c8/public