

## Product Name: UnderwriteAI

### Contacts:

- Product Manager: Zubair Akhtar
- Engineering Lead: TBD
- Design Lead: TBD

## 1. Executive Summary

This product requirement document describes building UnderwriteAI; an AI-powered credit decision support platform integrated with lending operations.

It responds to an identified market opportunity (reducing underwriting turnaround time from 45-60 minutes to 10-15 minutes while maintaining credit quality), is aligned with our strategy of scaling lending operations through intelligent automation, and leverages our competitive edge (domain expertise in SME and consumer lending) with Phase 1 MVP planned for Q2 2026.

### Success Criteria:

- Reduce average underwriting decision time from 50 minutes to 15 minutes (70% reduction)
- Increase underwriter productivity from 8 applications/day to 25 applications/day
- Improve decision consistency: reduce approval variance from 20% to <5%
- Maintain credit quality: default rates within  $\pm 2\%$  of baseline (current: 3.5%)
- Achieve 85%+ underwriter adoption within 3 months of launch
- Reduce 'verification status' support tickets by 30%

## 2. Market Opportunity

The global lending technology market has shown consistent growth of 12.4% CAGR (Deloitte, 2024) driven by digital transformation in financial services, regulatory pressure for faster credit decisions, and consumer demand for instant approvals. Currently in early growth stage, particularly in EU/UK markets where legacy systems dominate. TAM projected at €42 billion by 2028.

### Market Growth Stage:

Emerging to early growth. Digital lenders (Revolut, N26, Trade Republic) demand instant credit decisioning to compete with traditional banks while managing risk. SME and consumer lending segments show highest demand for automation.

### Current Pain Point:

Traditional underwriting processes create 45-60 minute delays per application. Internal data shows 31% drop-off during document submission phase for SME loans. Manual review creates inconsistent decisions (15-20% approval variance across underwriters) and operational bottlenecks requiring senior expertise.

## Potential Business Value:

Each 1% increase in underwriter productivity translates to approximately €240k in annual operational savings (based on current cost structure). Target 3x productivity improvement = €18M annual value. Additionally, faster turnaround improves customer conversion by estimated 15% (€3.2M incremental lending volume).

## 3. Strategic Alignment

UnderwriteAI aligns with our focus on becoming the leading provider of intelligent lending infrastructure for European financial institutions. It directly supports our mid-term objective of processing €500M in loan volume annually by removing the primary operational constraint (underwriter capacity) identified in 68% of partnership discussions (Q4 2025 Business Development retrospective).

### Vision Alignment:

Our vision is to make credit decisioning as fast and accurate as e-commerce checkout. Manual 45-minute reviews contradict 'fast'—this product resolves that tension while maintaining our commitment to responsible lending.

### Playing to Strengths:

We maintain our core differentiator (domain expertise in credit risk) while matching fintech competitors on decision speed. This is a 'table stakes' capability that unlocks our ability to scale, not a pivot from our risk-first approach.

### Company Objectives Supported:

- Short-term: Win 5+ enterprise lending partnerships in H1 2026 by demonstrating instant decisioning capability
- Mid-term: Scale loan processing capacity from 200 applications/day to 1,500/day without proportional headcount growth
- Long-term: Establish UnderwriteAI as the default choice for intelligent credit decisioning in Irish and UK lending markets

## 4. Customer & User Needs

### 4.1 Market Segments & Personas

#### Primary Segment: Mid-Level Credit Underwriters

- Examples: Underwriters at NBFs, digital banks, consumer finance companies
- Characteristics: 3-5 years experience, process 8-12 applications daily, pressure to balance speed with quality

#### Secondary Segment: Credit Managers & Risk Leads

- Examples: Heads of Credit, Risk Managers at lending institutions
- Characteristics: Manage 8-15 underwriters, responsible for portfolio quality, face scalability constraints

## User Personas:

### 1. "Mid-Level Underwriter Aoife" (Primary User - 70% of usage volume)

- Age: 28-35, processing SME and consumer loan applications
- Jobs-to-be-done: Review applications quickly, make confident credit decisions, meet daily quotas
- Pain points: 'Too much time on data entry, worried about making mistakes, inconsistent manager feedback'
- Constraints: Desktop-based workflow, needs to review 8-12 applications daily, limited time for complex analysis

### 2. "Senior Manager Conor" (Secondary User - 20% of usage volume)

- Age: 38-48, managing underwriting team of 10-15 people
- Jobs-to-be-done: Maintain consistent credit quality, identify training needs, scale team efficiency
- Pain points: 'High variance in underwriter decisions, training takes 3-6 months, can't scale team fast enough'
- Constraints: Budget limitations, regulatory compliance requirements, need audit trail

### 3. "Lending Operations Lead Siobhán" (B2B Stakeholder - 10% interaction)

- Role: VP Operations at lending institution
- Jobs-to-be-done: Reduce operational costs, improve processing capacity, ensure regulatory compliance
- Pain points: 'Manual underwriting doesn't scale, cost per loan is too high, need better data for risk management'
- Constraints: Limited budget for new technology, must integrate with existing loan origination system

## 4.2 Prioritized User Problems

### Problem 1: Manual Data Entry Consumes 30% of Underwriting Time (Severity: High, Frequency: Daily)

- Internal time-motion study shows underwriters spend average 15 minutes per application on data extraction from PDFs
- Importance: 9/10, Satisfaction: 2/10 (underwriter surveys, n=24 internal users)
- Impact: Lost productivity estimated at €420k annually across underwriting team

### Problem 2: Inconsistent Credit Decisions Create Risk Exposure (Severity: High, Frequency: Weekly)

- Approval variance of 15-20% across underwriters for similar risk profiles
- Senior review required for 40% of applications, creating bottlenecks
- Impact: Missed lending opportunities (overly conservative decisions) and potential credit losses (overly aggressive decisions)

### **Problem 3: Manual Review Creates Customer Friction (Severity: Medium, Frequency: Daily)**

- 45-60 minute turnaround time is uncompetitive vs. instant approval from fintech lenders
- 31% of SME applicants abandon during document submission phase
- Impact: Lost business to competitors offering instant decisions

#### **Geographic/Regulatory Constraints:**

- Phase 1 limited to Ireland and UK (GDPR-compliant infrastructure, FCA/CBI regulatory alignment)
- Must maintain compliance with Consumer Credit Act 1995 (Ireland), Consumer Credit Act 1974 (UK)
- AI credit decisioning requires explainability for regulatory audit (GDPR Article 22)

## **5. Value Proposition & Messaging**

### **5.1 Problem-Solution Mapping by Segment**

#### **For Mid-Level Underwriters (Segment 1):**

- Problem: 'I spend too much time on manual data entry and spreadsheet calculations'
- Solution: Automated document intelligence extracts financial data in <30 seconds
- Benefit: Process 25 applications/day instead of 8, spend time on edge cases instead of routine work
- Messaging: 'Focus on credit judgment, not data entry—AI handles the grunt work.'

#### **For Credit Managers (Segment 2):**

- Problem: 'Inconsistent underwriter decisions create risk exposure and audit issues'
- Solution: AI provides standardized risk scoring and policy compliance checks
- Benefit: <5% decision variance, full audit trail, reduce training time from 3 months to 1 month
- Messaging: 'Consistent credit quality at scale—every decision backed by data and policy.'

#### **For Operations Leaders (B2B):**

- Problem: 'Manual underwriting limits our lending capacity and increases cost per loan'
- Solution: 3x productivity improvement without proportional headcount increase
- Benefit: Scale from 200 to 1,500 applications/day, reduce cost per loan by 60%
- Messaging: 'Scale lending operations through intelligent automation—not just more people.'

### **5.2 Competitive Differentiation**

#### **vs. Manual Underwriting (Status Quo):**

- Similarity: Maintains human oversight and credit judgment
- Differentiation: 70% faster with AI-powered data extraction and risk scoring
- Message: 'Keep the expertise, lose the grunt work—AI assists, humans decide.'

### **vs. Fully Automated Credit Scoring (Fintech Competitors):**

- Similarity: Fast decision-making with ML-powered risk assessment
- Differentiation: Human-in-the-loop for complex cases and regulatory explainability
- Message: 'Speed of automation, judgment of experts—best of both worlds.'

### **vs. Legacy Lending Software (Temenos, FIS):**

- Similarity: Integrated with loan origination systems
- Differentiation: AI-native with RAG-based policy Q&A and document intelligence
- Message: 'Modern AI technology that integrates with your existing lending stack.'

## **5.3 Key Messaging**

**Headline: 'Underwrite loans in 10 minutes, not 50—with AI that explains every decision.'**

Supporting Points:

- '70% faster decisions: From 50 minutes to 15 minutes average processing time'
- 'Consistent credit quality: <5% variance across all underwriters'
- 'Full explainability: Every AI recommendation backed by policy citations and data'
- 'Human-in-the-loop: AI assists, underwriters decide—regulatory compliant and audit-ready'

## **6. Competitive Advantage**

### **6.1 Defensibility**

#### **What makes this defensible?**

Our competitive advantage stems from combining instant decisioning WITH human credit expertise—a combination pure automation can't match and traditional lenders can't achieve:

- Proprietary Risk Dataset: Our historical lending data (280k applications across SME, consumer, vehicle finance) trains a credit risk model competitors can't replicate. 12-month competitive lead.
- Domain Expertise: Deep knowledge of Irish/UK lending regulations, credit policies, and risk frameworks embedded in RAG knowledge base. Takes competitors 18+ months to build equivalent expertise.
- Integration Advantages: Pre-built connectors to major Irish/UK loan origination systems (Temenos, FIS). New entrants face 6-12 month integration cycles.
- Regulatory Compliance: GDPR Article 22 compliant explainability already built. Competitors entering EU market must build from scratch.

## 6.2 Why Competitors Won't Copy Immediately

### Pure Automation Players (Fintech Credit Scorers):

Their models optimize for speed, not explainability. Adding human-in-the-loop would slow their value prop. Unlikely to prioritize until regulatory pressure forces it (18+ month timeline).

### Legacy Lending Software (Temenos, FIS):

Enterprise software with 18-24 month release cycles. AI/ML capabilities require fundamental architecture changes. Minimum 2-year timeline for competitive feature parity.

### Consulting Firms Building Custom Solutions:

Custom builds cost €500k-€2M per implementation. Our SaaS model is 10x more cost-effective. Price advantage makes copying uneconomical.

## 6.3 Durability (3-Year Outlook)

Year 1: First-mover advantage in 'AI-assisted underwriting for Irish/UK lenders' positioning. Win deals competitors can't match on speed + explainability combination.

Year 2: Network effects from transaction data improve risk model accuracy faster than competitors. Portfolio performance data (default rates, loss given default) creates moat.

Year 3: Lending institutions have integrated UnderwriteAI into underwriter workflows and training programs. Switching costs prevent churn even if competitors achieve feature parity.

Risk Mitigation: Continue investing in model accuracy, expand to new lending verticals (auto, mortgage), and build partner lock-in through custom policy integration.

## 7. Product Scope and Use Cases

### 7.1 Key Capabilities & Features

#### Core Capability 1: AI-Powered Document Intelligence

Feature: Automated extraction of financial data from bank statements, tax returns, financial statements, and supporting documents

Tasks Product Must Perform Exceptionally Well:

- Accept PDF, Excel, image uploads via drag-and-drop (max 10 files, 25MB total)
- OCR text extraction from scanned documents with 95%+ accuracy
- Identify document type (bank statement, P&L, balance sheet, tax return, GST return)
- Extract key financial metrics: revenue, profit, cash flow, debt ratios, CIBIL score
- Calculate derived metrics: debt-to-income, DSCR, current ratio, ROA
- Complete extraction in <30 seconds for standard application (5-7 documents)

Desired Customer Outcome: Underwriters spend zero time on manual data entry; focus entirely on credit analysis and decision-making

Design Prototype: [<https://zubaira-create.github.io/underwriteAI-prototype/>] (UnderwriteAI Prototype)

High-Risk Assumption: OCR accuracy sufficient for financial documents with varied formats

Test Method: Pilot test with 500 real applications from 3 lending partners; measure extraction accuracy vs. manual review ground truth

## **Core Capability 2: RAG-Based Credit Analysis & Risk Scoring**

Feature: LLM-powered risk assessment using Retrieval Augmented Generation to analyze applications against credit policies and historical performance data

Tasks Product Must Perform Exceptionally Well:

- Generate credit risk score (300-900 scale) in <5 seconds after data extraction
- Provide score breakdown by component: Credit History (30%), Financial Health (35%), Business Stability (20%), Repayment Capacity (15%)
- Identify top 3-5 risk factors with severity ratings (High/Medium/Low)
- Cite specific policy sections that support or contradict approval
- Retrieve similar historical cases (3-5) with decision outcomes for comparison
- Maintain <2% hallucination rate (verified via human review)

Desired Customer Outcome: Underwriters make more consistent decisions faster; credit managers see <5% variance in approval decisions across team

Design Prototype: [[Link to Figma: Risk Analysis Dashboard v1.2](#)] (UnderwriteAI Prototype)

High-Risk Assumption: RAG retrieval accurate enough to cite correct policy sections without hallucination

Test Method: Shadow mode testing for 30 days comparing AI recommendations vs. senior underwriter decisions; measure alignment rate and policy citation accuracy

## **Core Capability 3: Natural Language Q&A for Policy & Risk Clarification**

Feature: Conversational AI interface allowing underwriters to ask questions about credit policies, risk factors, and calculation methodologies

Tasks Product Must Perform Exceptionally Well:

- Answer questions about risk scores: 'Why is the DSCR score low?'
- Explain policy compliance: 'Does this meet our debt-to-income threshold?'
- Provide calculation details: 'How is the current ratio calculated?'
- Compare with precedents: 'How have similar cases been decided?'
- Respond in <3 seconds with contextually relevant, policy-backed answers
- Maintain conversation context across 5+ exchanges

Desired Customer Outcome: Underwriters get instant answers to complex credit questions without escalating to senior staff; training time reduced from 3 months to 1 month

Design Prototype: [[Link to Figma: Q&A Assistant Interface v1.0](#)] (UnderwriteAI Prototype)

High-Risk Assumption: LLM can provide accurate, policy-compliant answers without requiring extensive fine-tuning

Test Method: Usability testing with 10 underwriters asking 50 common questions; measure answer accuracy and user satisfaction

## 7.2 User Flows

### User Flow 1: Aoife (Mid-Level Underwriter) - Standard SME Loan Review

- Aoife receives new €25,000 SME working capital loan application in queue
- Clicks 'Review Application' → sees applicant summary and document list
- Drags 5 documents (bank statements, tax returns, financials) into upload area (30 seconds)
- AI extracts financial data, calculates ratios, generates risk score (30 seconds)
- Aoife reviews risk score (650 - Medium Risk), sees key metrics and risk factors (3 minutes)
- Asks AI: 'Why is the debt-to-income ratio flagged?' → Gets instant explanation (30 seconds)
- Reviews AI loan recommendation: €20-22k amount, 36 month tenure, 12.5% interest rate (2 minutes)
- Aoife agrees with recommendation, clicks 'Approve' with €22k amount (1 minute)
- Decision logged with full audit trail; applicant notified (automated)

**Total time: 10-12 minutes (vs. 50 minutes baseline)**

### User Flow 2: Conor (Senior Manager) - Override Review

- Conor receives alert: Aoife overrode AI recommendation (approved €25k instead of €22k)
- Opens application in review dashboard, sees AI analysis and Aoife's decision reasoning
- Reviews risk factors, asks AI: 'What's the default rate for similar profiles?' (1 minute)
- Sees 3 similar historical cases: 2 performed well, 1 defaulted after 18 months
- Agrees with Aoife's judgment, approves override (2 minutes)
- Sends feedback to Aoife: 'Good catch on the seasonal cash flow pattern' (1 minute)

**Total time: 4 minutes (vs. 25 minutes baseline for manual review)**

## 8. Non-Functional Requirements

### 8.1 General Requirements

#### Performance:

- Document processing: <30 seconds for standard application (5-7 documents)
- Risk score generation: <5 seconds after data extraction
- Q&A response time: <3 seconds for typical queries
- Dashboard load time: <2 seconds for underwriter queue
- System availability: 99.5% uptime during business hours (9 AM - 6 PM GMT)

#### Scalability:

- Support 100 concurrent underwriters
- Process 500 applications/day initially, scale to 2,000/day within 6 months
- Handle 50,000 stored applications in first year without performance degradation

- Support multi-tenant architecture for future white-label opportunities

#### **Reliability:**

- 99.5% uptime SLA for core services (max 3.6 hours downtime/month)
- Graceful degradation: if AI services fail, route to manual review workflow
- Auto-retry logic for transient failures (3 attempts with exponential backoff)
- Data backup: hourly incremental, daily full backup with 7-year retention

#### **Security:**

- GDPR compliance: Data encryption at rest (AES-256) and in transit (TLS 1.3)
- Role-based access control (underwriter, manager, admin, auditor roles)
- Immutable audit logs for all credit decisions and AI recommendations
- PII masking in logs and analytics dashboards
- Compliance: FCA (UK), CBI (Ireland), GDPR, ISO 27001 alignment

#### **Peak Load Expectations:**

- Month-end processing: 1,000 applications/day across all clients
- New partner onboarding: 2,000 applications in first week for single large lender

## **8.2 AI-Specific Requirements**

#### **AI Architecture:**

- Risk Scoring: Gradient boosting classifier (XGBoost) trained on historical loan performance data
- Document Intelligence: Azure Document Intelligence + custom OCR post-processing
- RAG System: Vector database (Pinecone/Weaviate) with OpenAI embeddings for policy retrieval
- LLM: GPT-4 or Claude 3 Opus for risk analysis and Q&A
- Fallback Logic: Rule-based scoring if ML model unavailable (conservative thresholds)

#### **Accuracy & Reliability:**

- OCR Accuracy:  $\geq 95\%$  character recognition for financial documents
- Data Extraction Accuracy:  $\geq 90\%$  for structured fields (revenue, profit, ratios)
- Risk Score Consistency:  $< 5\%$  variance on same application reprocessed
- LLM Hallucination Rate:  $< 2\%$  (measured via policy citation accuracy)
- Credit Quality: Default rates within  $\pm 2\%$  of baseline after AI implementation

#### **Ethical Standards:**

- No discrimination based on protected characteristics (GDPR Article 22, Equal Status Acts)
- Bias testing: Model performance variance  $< 5\%$  across demographic groups (age, gender, geography)
- Explainability: Every AI decision must include reason codes and policy citations
- Human oversight: All high-risk applications (score  $< 600$ ) require manual review

### **Measurement & Validation:**

- Primary Metric: Precision and Recall for default prediction (target: F1 Score 0.85+)
- Secondary Metrics: Approval consistency (<5% variance), processing time (target: 15 min avg)
- A/B Testing: Shadow mode for 30 days comparing AI vs. manual decisions before full rollout
- Ongoing Monitoring: Weekly model performance reports; monthly retraining evaluation

### **Maintaining Quality Over Time:**

- Automated feedback loop: Actual default outcomes fed back into training data quarterly
- Quarterly model audits by Risk & Compliance team to validate bias metrics and accuracy
- Lending segment-specific tuning: Different risk thresholds for SME vs. consumer vs. vehicle loans
- Drift detection: Alert if approval rate changes >10% or default rate exceeds baseline by >2%

## **9. Go-to-Market Approach**

### **Phase 1: MVP Pilot (Weeks 1-12, Target: Q2 2026)**

#### **Target Segment: Internal underwriting team (24 underwriters) at pilot lending partner**

##### **Build Scope:**

- Document intelligence with OCR and financial data extraction
- Risk scoring model (basic version with 4 risk components)
- Underwriter review dashboard (queue management, application details, decision buttons)
- Basic audit trail and reporting
- Integration with pilot partner's LOS (API-based data exchange)

##### **Release Strategy:**

- Pilot Partner: 1 Irish consumer finance lender (€50M annual loan volume)
- Traffic Rollout: 20% of new applications routed to UnderwriteAI, 80% to legacy process (A/B test)
- Duration: 12 weeks with weekly progress reviews

##### **Success Metrics:**

- 50% reduction in processing time (target: 25 minutes from 50 minutes baseline)
- 85% data extraction accuracy on pilot application set (n=500 applications)
- Underwriter satisfaction: 7/10 or higher (survey after 12 weeks)
- Credit quality maintained: Default rate within  $\pm 2\%$  of control group
- Pilot partner commits to Phase 2 expansion

#### Value Demonstration:

- Weekly reports showing time savings and consistency improvements
- Case study: 'Consumer Finance Lender Processes 40% More Applications with Same Team'
- ROI calculation: €420k annual operational savings projection
- 

#### **Phase 2: Enhanced Intelligence & Scale (Weeks 13-26, Target: Q3 2026)**

#### Expansion Scope:

- RAG-based Q&A assistant for policy questions and risk clarification
- Loan structure recommendations (amount, tenure, pricing)
- Similar case retrieval for precedent comparison
- Manager override dashboard with approval analytics
- Expand document types: GST returns, balance sheets, cash flow statements

Target Segment: Add 2 new lending partners (1 SME lender, 1 vehicle finance company)

#### Release Strategy:

- 100% traffic to UnderwriteAI for pilot partner (graduate from A/B test)
- Onboard 2 new partners with full feature set
- Geographic expansion: UK launch with localized compliance

#### Success Metrics:

- 70% time reduction achieved (15 minutes average processing time)
- Decision consistency: <5% approval variance across underwriters
- Q&A accuracy: 90% of responses rated as 'helpful' by underwriters
- 3 active lending partners using UnderwriteAI
- 500+ applications processed daily across all partners

#### Win Conditions:

- Win 2 competitive deals vs. legacy lending software citing AI capabilities
- Partner NPS >50 (baseline: 32 for manual underwriting)
- Feature coverage: 85% of underwriter workflows supported

### **Phase 3: Market Expansion & Optimization (Weeks 27-52, Target: Q4 2026 - Q1 2027)**

#### Advanced Features:

- Enhanced ML risk model v2.0 with behavioral signals and transaction velocity analysis
- Real-time portfolio monitoring dashboard for credit managers
- Custom policy configuration for each lending partner
- Multi-language support (Irish, Polish, Romanian for diverse underwriter teams)
- Mobile app for field underwriters (vehicle inspection photos + credit analysis)

Target Segment: All SME, consumer, and vehicle finance lenders in Ireland and UK

#### Success Metrics:

- 10+ lending partners actively using UnderwriteAI
- €500M+ annual loan volume processed through platform
- €18M operational savings delivered across customer base
- Default rates maintained within  $\pm 2\%$  of pre-AI baseline
- Market position: Recognized as leading AI underwriting platform in Ireland/UK

#### Market Position:

- Case studies with 5+ partners demonstrating measurable ROI
- Thought leadership: Conference presentations at Fintech Connect, Money20/20 Europe
- Competitive win rate  $>70\%$  in RFPs where AI decisioning is requirement

## Appendix A: Open Questions & Decisions Needed

- Risk Tolerance Decision (Owner: CFO + Head of Risk, Deadline: Week 2)
  - What default rate increase (if any) are we willing to accept for 70% processing time improvement?
  - Recommendation: Maintain default rate within  $\pm 2\%$  of baseline; prioritize speed without sacrificing credit quality
- Human Override Policy (Owner: Product + Risk, Deadline: Week 3)
  - Should underwriters be allowed to override AI recommendations without manager approval?
  - Recommendation: Allow overrides but require written justification; manager review for  $>€20k$  overrides
- Pricing Strategy (Owner: Finance + Product, Deadline: Week 4)
  - Per-application pricing vs. platform licensing vs. savings-based pricing?
  - Recommendation: Platform licensing (€8k/month per 10 underwriters) to maximize adoption and predictable revenue
- LLM Provider Selection (Owner: Engineering + Procurement, Deadline: Week 5)
  - OpenAI GPT-4 vs. Anthropic Claude vs. open-source LLM for RAG system?
  - Recommendation: Claude 3 Opus for lower hallucination rate and better policy citation accuracy in initial testing
- Explainability Standards (Owner: Legal + Compliance + Product, Deadline: Week 6)
  - What level of AI decision explainability satisfies GDPR Article 22 and CBI regulatory expectations?
  - Recommendation: Every AI score must include: reason codes, policy citations, similar case comparisons, manual override capability

## Appendix B: Success Metrics Dashboard (Post-Launch Monitoring)

### Weekly Tracking:

- Average processing time per application (target: 15 minutes)
- Data extraction accuracy rate (target: 90%+)
- Approval variance across underwriters (target: <5%)
- AI recommendation override rate (baseline: track for learning)
- System uptime and API response times

### Monthly Tracking:

- Underwriter productivity: applications processed per day per underwriter
- Credit quality: default rate vs. pre-AI baseline
- Partner satisfaction: NPS surveys
- Model performance: precision, recall, F1 score for default prediction
- Cost per application: operational savings delivered

### Quarterly Business Review:

- Total operational savings delivered (target: €18M annually at scale)
- Partner retention rate and expansion opportunities
- Competitive win/loss analysis (deals where AI was deciding factor)
- Product roadmap prioritization based on partner feedback
- Regulatory compliance audit results

## Appendix C: Competitive Analysis

Competitor	AI Decisioning	Processing Time	Explainability	Target Market
Manual Underwriting (Status Quo)	✗ No	45-60 minutes	☑ Full (human judgment)	All traditional lenders
Pure Automation (Fintech Scorers)	☑ Yes	< 1 minute	⚠ Limited (black box)	Digital lenders, fintechs
Legacy Lending Software (Temenos, FIS)	✗ No	30-45 minutes	☑ Full (rule-based)	Enterprise banks
UnderwriteAI	☑ Yes (Human-in-loop)	10-15 minutes	☑ Full (RAG-based)	Irish/UK SME & consumer lenders

### Key Competitive Insight:

Pure automation (fintech scorers) achieves speed but lacks explainability for regulated lenders. Legacy software provides explainability but is too slow. UnderwriteAI uniquely combines speed + explainability through human-in-the-loop AI, positioning us in the 'sweet spot' for Irish/UK lending institutions.

### Document Status: Draft v1.0 for stakeholder review

#### Next Steps:

- Schedule technical feasibility review with Engineering (Week 1)
- Legal/Compliance sign-off on AI decision framework (Week 2)
- Identify pilot lending partner (Week 3)
- Finalize success criteria for pilot → full rollout decision (Week 4)

### Prototype Link:

<https://zubaira-create.github.io/underwriteAI-prototype/>