

Risk Register

3-month feasibility study (Oct - Dec 2026) | Total cost: GBP 50,000

ID	Category	Risk Description	Likelihood	Impact	Mitigation Strategy	WP	Status
R1	Technical	Convergence proofs do not hold at 100K+ particles in 384 dimensions	Medium	High	Identify sufficient conditions (bounded insertion rate, minimum decay constant) for convergence. Document operational envelope. Partial convergence results remain valuable for Phase 2 planning.	WP1	Open
R2	Technical	Latency overhead exceeds 500ms at 100K entries, making Hippo impractical	Medium	High	Benchmark includes latency at every scale point. If exceeded, identify optimisation strategies: spatial indexing, approximate force calculations, GPU acceleration via RTX 5080. Document in Phase 2 spec.	WP2	Open
R3	Technical	No advantage over baselines on standard retrieval metrics (MRR, NDCG)	Low	Medium	Benchmark matrix includes both standard metrics AND Hippo-unique capabilities (consolidation, forgetting, degradation resistance). Parity on standard metrics plus unique capabilities is a strong commercial result.	WP2	Open
R4	Technical	Multi-agent shared memory introduces unresolvable conflict dynamics	Medium	Medium	WP3 conflict resolution simulation will identify failure modes early. Fallback: specify partitioned memory spaces with controlled merge points rather than fully shared particle space.	WP3	Open
R5	Commercial	Insufficient interest from Phase 2 consortium partners	Low	Medium	Partner engagement begins Month 1, not deferred to Month 3. Convergence proofs and benchmark results provide concrete evidence for recruitment. Target multiple academic groups and enterprise contacts in parallel.	WP3	Open
R6	Managerial	Key-person dependency (single researcher for 3-month project)	Low	High	All work version-controlled and documented. 527-test regression suite. Short project duration limits exposure. Weekly progress snapshots ensure continuity if disruption occurs.	WP4	Mitigated
R7	Technical	Evaluation corpus not representative of real agent workloads	Low	Medium	Use 10 diverse synthetic workload profiles modelled on real use cases (customer support, legal, clinical, research, manufacturing). Supplement with publicly available datasets from multiple domains.	WP2	Open
R8	Financial	Cloud compute costs exceed budget during intensive benchmarking	Low	Low	Primary compute on local RTX 5080 GPU (zero marginal cost). Cloud used only for reproducibility verification. Usage alerts set at 80% of cloud budget allocation (GBP 5,000). Total project contingency available.	WP2	Mitigated

Risk Assessment Matrix

	Low Impact	Medium Impact	High Impact
Medium		R3, R4	R1, R2
Low	R8	R5, R7	R6

Risk review: weekly during project execution. Principal researcher (Keith So) has overall risk ownership.

Risks reviewed at each milestone gate (M1-M4). Status updated as mitigations are implemented.