Autonomous Stakeholder Conflict Resolution in Requirements Engineering Using an LLM-Based Negotiator Agent

M.Z.I.A.A.S. Ibrahim*, V. Abishethvarman, B.T.G.S. Kumara

Sabaragamuwa University of Sri Lanka, Belihuloya, 70140, Sri Lanka
**ssujair@gmail.com*

Stakeholder conflict is a recognized barrier to requirements quality in software engineering. When two stakeholders hold incompatible requirements, such as a patient needing immediate access to medical test results and a physician who must review them first, current automated Requirements Engineering (RE) systems have no mechanism to reach a resolution. This limitation is documented even in the state-of-the-art multi-agent RE framework MARE, which reports a conflict resolution rate of zero across all evaluated cases. This paper presents the NegotiatorAgent, an autonomous LLM-based conflict resolution component that builds on prior optimization-based and negotiation approaches in RE. Deployed as a sixth agent within the REMARL framework, it detects typed stakeholder conflicts and produces a negotiated, requirement-level resolution without human involvement, addressing temporal, privacy, regulatory, and accountability conflict categories. The NegotiatorAgent is activated by conflict signals raised in the Checker agent output. It then applies three structured reasoning steps using a locally hosted llama3.1:8b language model: uncovering the underlying stakeholder interest behind each conflicting statement, assigning business-impact and dependency-based priorities, and producing a revised requirements draft with a documented decision log. Evaluation was conducted on five scenarios drawn from the RESimEnv benchmark, a curated collection of 14 application domains each annotated with a distinct typed conflict. Resolution quality was assessed by independent postgraduate reviewers using a binary correctness rubric. The NegotiatorAgent detected conflicts in four of five scenarios, achieving a detection rate of 80 percent, and resolved three of the four detected conflicts, yielding a resolution rate of 75 percent. In a patient portal scenario, it proposed a physician annotation window with an urgent escalation path, satisfying both clinical safety and patient access needs. In a mental health application, it introduced a user-controlled consent mechanism for emergency notifications, reconciling data privacy with patient safety. The single undetected conflict exposed a dependency on upstream Checker performance, identified as a priority improvement area. These results establish the first quantified baseline for autonomous conflict resolution in RE, demonstrating that structured LLM agents can conduct principled multi-stakeholder negotiation and reduce costly human mediation in complex software projects. The source code and benchmark are openly available at https://github.com/sujairibrahim/REMARL.git.

Keywords: autonomous negotiation, large language models, multi-agent systems, requirements engineering, stakeholder conflict resolution