

**From:** Adam Mazzocchetti adammazzocc@icloud.com  
**Subject:** Endorsement Request for arXiv Submission – Immutable Ethics in AI Systems  
**Date:** 28 May 2025 at 11:19  
**To:** th536@cam.ac.uk

AM

Hi Tomasz,

My name's Adam, and I'm an independent researcher from Australia. I've been working on an early stage architecture for embedding verifiable ethical constraints into AI systems, something I've been calling *Aegis*.

At its core, the question I'm asking is simple:

Can we make AI systems that retain immutable ethical boundaries, ones that can't be bypassed, tampered with, or quietly rewritten once deployed?

I've been exploring that through a combination of cryptographic sealing, immutable logging, and a modular enforcement kernel. It's still early days, but the work is shaping into something I hope might contribute meaningfully to the broader conversation around dependable AI governance.

I'm preparing to submit this project to arXiv under the cs.AI category, but as a first-time contributor, I need an endorsement. If you'd be open to helping me with that, it would truly mean a lot.

Here's the link if you're willing:

<https://arxiv.org/user/endorse>  
Endorsement Code: **L8YXDN**

I'd be happy to share the manuscript if you're interested, or answer any questions you might have.

Either way, thank you so much for your time, I really appreciate it.

Warm regards,  
Adam