

Chapter 17 Checklist: AI Usage Safety Protocol for Regulatory Work

Subtract to Ship: MDR Certification*, Chapter 17: The AI Advantage

Instructions

If you are using AI for any aspect of regulatory work, run this protocol. AI saves 80% of time on structured tasks but hallucinates in ways that look perfectly credible. This checklist ensures AI output is verified before it enters your regulatory documentation. The rule: AI flags, the expert adjudicates. Not the other way around.

Part 1: AI Usage Map

Document where AI is used in your regulatory workflow:

Task	AI Used?	AI Tool	Human Reviewer	Review Depth
Document drafting (QMS procedures)	Yes No			Full Spot-check
Literature research / screening	Yes No			Full Spot-check
Gap analysis (tech doc vs. requirements)	Yes No			Full Spot-check
Safety database monitoring (PMS)	Yes No			Full Spot-check
Risk analysis drafting	Yes No			Full Spot-check
Clinical evaluation drafting	Yes No			Full Spot-check
Template generation	Yes No			Full Spot-check
Classification decisions	Yes No			Must be FULL
Other: _____	Yes No			

Part 2: Where AI Helps vs. Where AI Is Dangerous

Safe for AI (with human review)

- ☐ Document drafting (first drafts, structure, formatting)
- ☐ Literature research (scanning, filtering, extracting key data)
- ☐ Gap analysis (structural gaps in documentation)
- ☐ Safety database pre-categorisation
- ☐ Template generation (QMS procedures, SOPs, forms)
- ☐ Cross-reference checking

Dangerous with AI (requires expert judgment, not pattern matching)

- ☐ Classification decisions (Annex VIII rule interpretation)
- ☐ Clinical evidence sufficiency assessment
- ☐ Risk analysis (identifying failure modes requires clinical imagination)
- ☐ MDR article interpretation (ambiguous articles, article interactions)
- ☐ Intended purpose definition
- ☐ Benefit-risk determination

If AI is being used for any "dangerous" task: ensure the output is treated as a starting point only, with full expert review.

Part 3: The One-Paragraph Verification Test

Pick one AI-generated paragraph from your regulatory documentation. Right now.

Step	Result
Paragraph source: (document, section)	
Time to verify:	____ minutes
MDR article citations verified?	All correct ____ errors found
Standard references verified (correct edition)?	All correct ____ errors found
MDCG guidance references verified (exist and current)?	All correct ____ errors found
Factual claims traceable to specific MDR provisions?	All correct ____ errors found
Total errors found:	____

Assessment

Errors	Interpretation
0	Good, but one paragraph is a small sample. Run on more.
1-2	Multiply across a 30-page document = dozens of potential problems.
3+	Systemic issue. AI output cannot be trusted without line-by-line verification.

Part 4: Anti-Complacency Protocol

After AI gets 10 results right in a row, reviewers stop checking. Build these structural safeguards:

Verification Techniques

Technique	Implemented?
Adversarial verification: Reviewer's job is to FIND errors, not confirm correctness. "Find three errors in this output."	Yes No
Spot-check rotation: Random subset of AI outputs selected for deep verification against actual MDR text.	Yes No
Two-human review for critical documents: Risk analyses, clinical evaluations, classification rationale, never AI-generated and single-reviewed.	Yes No
Source verification: Every MDR article, standard, and MDCG reference in AI output checked against actual text.	Yes No

Part 5: AI-Generated Document Approval Gate

Before any AI-assisted document enters the technical file or NB submission:

Check	Passed?
Every MDR article citation verified against current consolidated text	
Every standard reference verified as current edition	
Every MDCG guidance reference verified as latest version	
Every factual claim about MDR requirements traced to specific article/annex	
No hallucinated references (non-existent articles, standards, guidance)	
Content reviewed by someone qualified to catch domain-specific errors	
Two independent reviewers for critical documents (risk, clinical, classification)	

Part 6: Qualified Reviewer Check

Question	Answer
Who on your team can catch AI errors in regulatory documentation?	
Do they have domain expertise to identify hallucinated article references?	Yes No
Can they spot misapplied classification rules?	Yes No
Can they identify clinically implausible risk ratings?	Yes No

If nobody is qualified to catch AI errors: AI is not making your work faster. It is accumulating undetected errors.

Part 7: EU AI Act Check (if your device includes AI)

Question	Answer
Does your device include AI/ML components?	Yes No
If yes: Is the AI a safety component of the medical device?	Yes No Unsure
If yes: Is the device classified as high-risk under AI Act Annex I?	Yes No Unsure
Have you assessed additional AI Act requirements (transparency, data governance, human oversight)?	Yes No

If your device includes AI: flag this for your regulatory partner. MDR + AI Act interaction is still being clarified.

Part 8: AI Briefing Protocol (run BEFORE you ask)

Model-agnostic. Works with any capable AI assistant, now and after the tools change.

Before the question

Briefing step	Done?
Context handed over: device description, intended purpose statement, classification rationale	
Relevant filled worksheet(s) from this playbook included	
Request is specific and checkable (draft structure, gap list, counter-arguments), not "write my X"	
Instruction added: "Name the specific MDR article, annex, or guidance for every regulatory claim"	

Judging the answer

Good answer	Bad answer
Cites specific provisions you can look up	Cites nothing you can verify
Separates "the regulation states" from interpretation	Blends regulation and opinion into one confident voice
Flags what it is unsure about	Sounds certain about everything
Asks about your device	Assumes facts about your device

Push-back (mandatory for anything that matters)

Push-back move	Asked?
"What did you assume about my device that I did not tell you?"	
"Argue against your own recommendation."	
Re-ask with one constraint changed; check for quiet contradictions	

If two answers contradict each other: you have reached the edge of pattern matching. Stop. Bring in a human expert.

Every output from this protocol is a DRAFT until it passes the Part 5 approval gate.

Summary Rule

AI is an accelerator, not an autopilot.

The expert in the loop is not optional. The expert is the entire point.

- ☐ I understand that "the AI wrote it and it looked correct" is not a defence before an NB or competent authority
- ☐ Article 19(3) makes the manufacturer responsible for compliance, not the AI tool
- ☐ Every AI output in our regulatory documentation has been verified by a qualified human

Date completed:

