An Integrated Maturity Model for modern software development: addressing Security, Artificial Intelligence and Governance in DevOps

Nuno Seixas[1][0000-0002-3744-5746], Paulo Rupino da Cunha[1][0000-0003-2701-5248],
Marco Vieira[1][0000-0001-5103-8541]

[1] University of Coimbra
{naseixas,rupino,mvieira}@dei.uc.pt

Abstract

Software is now a strategic asset for almost every business. As such, new technologies and features are needed, and so, AI-infused solutions are becoming more frequent. Also, as software is part of every infrastructure, even critical ones, it becomes a more attractive target for cybersecurity attacks. To build software that can be seen as trustworthy, organizations need to improve their processes and integrate different skills, methods, capabilities, and practices. A way to achieve it is by using a maturity model that can lead that improvement path. Since the current maturity models present limitations – our literature review shows they are considered to be static, focused on specific contexts, and built with a closed architecture, we set ourselves with the goal of building a new maturity model that addresses these limitations. Our goal is to develop a model that can integrate capabilities and practices useful for producing AI-infused software with cybersecurity in mind, based on an open-architecture that allows for integrating future domains and be evolved by organizations that use it. We defined an iterative methodology for producing this maturity model that integrates in every iteration a validation step by academic and industrial partners. As we finish the first iteration, which addresses cybersecurity capabilities, we are starting the validation process for this increment, but preliminary results show our methodology to be leading to sound results.

Keywords: Maturity Model, Cybersecurity, Artificial Intelligence, DevOps, Integrated

Acknowledgments. This work is funded by the FCT - Foundation for Science and Technology, I.P./MCTES through national funds (PIDDAC), within the scope of CISUC R&D Unit - UIDB/00326/2020 or project code UIDP/00326/2020

1 Introduction

Software has become a central part of almost any successful business, being considered an essential tool for organizations to operate effectively and efficiently, allowing companies to increase productivity and enhance customer experiences [1].

The advent of Artificial Intelligence (AI) is making software even more critical for supporting businesses. AI algorithms can automate repetitive tasks, analyze large amounts of data to identify trends and patterns, and make predictions that can help businesses make better decisions [2]. However, with increased reliance on software and AI comes an increased risk of cybersecurity threats. As businesses collect and store more data, they become a more attractive target for cybercriminals [3]. For software to produce its value, it needs to be reliable, consistent, and transparent. Therefore, these two trends need to be built on top of a strong foundation – using DevOps as the underlying set of principles for software development [4].

Looking at the challenges posed by AI and cybersecurity over DevOps, our initial research showed that a new, integrated software development approach is needed. Any organization that produces software will need to consider the necessary skills, methods, approaches as an interconnected network of capabilities and not as two different isolated sets of practices. Only then can those organizations maximize the probability of success while decreasing the risks of failure when infusing their software with AI, as stated in [2]. For this purpose, we defined as the main objective of this work the definition of a maturity model that, by definition, can help understand the initial status, foresee the end status, and plan the path from one to the other. Additionally, this integrated maturity model will need to be able to address the interconnected set of capabilities that arise from applying both AI techniques and cybersecurity principles over a DevOps foundation. This proposed model also needs to address some of the critiques that have been appointed to maturity models, such as being static, linear, and focused on a single definition of the right performance and improvement path [4, 5]. Therefore, this new maturity model will need to be able to adapt to different business contexts but also to be able to integrate different dimensions that proved important for the software engineering organizations.

As for the organization of this article, in section 2 we discuss the state-of-the-art on topics related to this work, namely focusing on the existing maturity models for AI and for cybersecurity. In section 3 we present research objectives that guide our work and the methodology used. In section 4 we present the work that we did, and some obtained preliminary results. In section 5 we present the planned work and expected results. Finally, in section 6, we present the conclusions for the presented work.

2 State-of-the-art

DevOps as a foundation, brings speed and quality to the software releases [4, 6]. The current business context demands shorter feedback loops in producing, deploying, and operating software. But for these loops to be rightly executed, a software release needs to have quality and be consistent [4]. As for Governance, there are a set of different frameworks that can be used as reference for IT related organizations, each one providing different approaches but with one same goal: to achieve excellence and trust [7]. Nonetheless, the most well-known frameworks such as ITIL [8], COBIT [9] or CMMI [10] do not explicitly address the interdependency of using AI and cybersecurity.

In general, maturity is considered as a state of being complete, perfect, or ready [11]. This definition is aligned with a biological perspective that presents something as mature when it has reached its full development capacity, as a final state. But from a psychological perspective, Blank et al [12] present maturity as “the extent to which a person acts independently, is able and willing to take responsibility, and desires to achieve”. The same authors consider this perspective more adequate to an organizational context, as organizations are built based on human relationships, and, therefore, on psychological and sociological interactions. As such, in this work, we will adopt their view, presenting organizational maturity as the capacity to understand its context, and take actions to adapt to it, therefore maximizing the probability of reaching their goals.

One of the most well-known maturity models in Information Technology (IT) [13], is the Software Engineering Institute (SEI)’s Capability Maturity Model (CMM) and its evolution, CMMI [10], currently in its version 3.0, which has been used to guide and evaluate different paths of improvement taken by different organizations. Although maturity models have historically been considered as an important tool for process improvement [13], there is a growing critique around them [4, 5]. They are considered static, based on one single end-goal and a unique improvement path. All those discussions on why maturity models are failing, confirm the need for an improvement tool that is context-aware, multidimensional, and dynamic.

To understand the state-of-the-art for maturity models that could connect cybersecurity and AI on top of DevOps and Governance, we conducted a multivocal systematic literature review [14] to answer the following research question:

RQ: What is the current state of the art on software development maturity models that address cybersecurity and the infusion of AI in software?

We followed the guidelines proposed by Kitchenham [15] on how to perform a Systematic Literature Review (SLR) in the computer science domain. Furthermore, following the recommendations by Garousi et al. [14], we also integrated in our multivocal systematic literature results grey literature (e.g., technical reports, books, and blog posts). We then executed the SLR following Kitchenham [15]’s guidelines and Adams et al. [16]'s recommendations. As step 1, the following search expressions were defined:

1. (MLOps or AI or AIOPS or “Cognitive Systems”) AND “maturity model”
2. (DevSecOps OR SecDevOps OR Security) AND “maturity model”

The search (step 2) was executed between January and May 2022 on the EBSCOhost databases (Academic Search Complete, Business Source Complete, and EconLit with Full Text) and considering only documents written in English, Spanish or Portuguese. In steps 3 and 4, we did an initial triage over the initial 1783 results, eliminated repeated ones, leaving us with1552 results. In step 5, we screened the results by reading the title and abstract to identify papers that were good candidates for full-text reading. This resulted in 125 works as several were discarded for two reasons; (i) results associated with life sciences' studies that use the term "maturity model"; and (ii) results written in a language other than English, Spanish or Portuguese. Finally, in step 6, we analyzed the full text of all those results.

None of the remaining results presents maturity models for software development that simultaneously address the concerns related to cybersecurity and AI infusion. Our findings fall into one of two groups: Cybersecurity Maturity Models and Artificial Intelligence Maturity Models, which we present in the next paragraphs.

Cybersecurity maturity models

Different authors present definitions for cybersecurity, but in this work we'll be using the one from Hoang and Le [17], who states that "*cyber security can be considered as a collection of systems, tools, processes, practices, concepts and strategies that are used to prevent and protect the cyber space from unintended interaction and unauthorized access and to preserve the confidentiality, integrity, availability, authenticity, accountability (CIAAA) and other properties of the space and its resources*". This definition accounts for different factors – systems, tools, processes, practices, concepts, and strategies – to achieve not just one specific gain but an interconnected set of gains over confidentiality, integrity, availability, authenticity, and accountability.

The literature shows that cybersecurity is an important topic because of two factors. First, a growing business complexity, with high degree of volatility, and use of heterogeneous systems [3]. Second, software is now part of critical infrastructures, and these are increasingly vulnerable to cyber-attacks, as stated in [18].

A relevant conclusion for our work stems from Frijns et al. [3], who states that, although cybersecurity is important and vulnerabilities represent a high risk of disrupting business, current software engineering methods still do not address it explicitly. In fact, looking at different Agile methods and DevOps, the same author confirms there is not an explicit call to security practices or safeguards that address the increasing risk of cybersecurity problems. Also, different works confirm that there is a need for a new Maturity Model that explicitly addresses both DevOps and Security, therefore, being able to handle a more complete set of threats [19, 20].

In the SLR we identified two groups of security maturity models: one focused on organizational practices that can be used by any technology-related organization, and another focused on the software development lifecycle. In the first group, we find examples such as C2M2 [21], NICE [22], NIST Cybersecurity Framework [23] and CMMC [24]. In the second group, focused on the software development lifecycle, we identified two instances: BSIMM [25] and SAMM [26].

After analyzing the literature related to cybersecurity, we have identified some remaining challenges. First, the existing models, focused on organizational practices, do not address how to include them into the software development lifecycle. Second, confirming the critiques to maturity models, all those that we identified have a closed architecture and do not allow for integrating additional perspectives or contexts. Third, there is a lack of a common language, making it very difficult to integrate different models, with different levels and different quantification schemas.

AI maturity models

The SLR results show that there is no maturity model focused on organizations implementing AI methods or more specifically, Machine Learning. There are some preliminary proposals from Alsheibani et al.[27], Desouza et al.[28] and Akkiraju et al.[29] where the authors confirm the need for defining an organizational capability to ensure the return on investment (ROI) that AI promises. We further identified three main trends related to AI: (i) Governance and Compliance; (ii) Transparency and Explainability, and (iii) MLOps.

As for Governance and Compliance, we confirm an increasing demand for regulation, especially in industries like finance and healthcare, as confirmed by Candelon et al. [30]. These authors consider regulation as essential for the development of AI systems, not as a limit to its capabilities but to make consumers trust them. This regulation work is being currently led by the European Union (EU) [31] with its 2021 proposal for an AI legal framework based on a risk assessment.

Regarding Transparency and Explainability in AI, different authors talk about the need for these systems to be able to explain their own rational and, ultimately, prove themselves as trustworthy, as explicitly said by Mora-Cantallops et al. [32], for whom transparency is one of the key requirements for trustworthy AI.

Finally, regarding MLOps, and as mentioned by Dang et al. [33], the advent of cloud computing changed the paradigm of producing and releasing software in the last decades. In this context, software releases moved from a traditional boxed product to a set of services being continuously released. AI-infused products, like any other kind of software, need also to be able to be deployed with a consistent and trustworthy approach.

Analyzing the results from the SLR, we find that, as for cybersecurity, there are some challenges left answered. First, the existent work doesn’t address how an organization should integrate different skills and mindsets to get the promised return on investment. Second, there isn’t a definition of the kind of governance that is needed, namely, by the implementation of regulation towards explainability, transparency and risk analysis. Third, from an operational perspective, it is unclear what kind of software development methods are needed to achieve consistency in development and release of AI-infused software artifacts.

3 Research objectives and methodological approach

Taking into consideration the current state-of-the-art, we defined as objective for our work to produce an integrated maturity model that addresses Cybersecurity and AI, over a foundation of DevOps and Governance.

The proposed maturity model needs to be adaptable to different business contexts, for organizations producing software that has needs related to the two mentioned domains. This maturity model should be sustainable, i.e., it must have an open architecture, built in such a way that further domains can be integrated without having to change the model structure.

We are following a Design Science Research (DSR) approach, as defined by Peffers and Hevner [34, 35], complemented with a methodology for building maturity models presented by Becker [13], resulting in an iterative approach, as presented in **Fig. 1**.

Step 1 is meant to produce an identification and characterization of the problem, which is in part achieved by the systematic literature review explained in the previous section, but also with a comparison of the existent maturity models for each one of the two domains – Cybersecurity and AI. Step 2 is meant to produce a definition of objectives and goals for the proposed maturity model, but also, to produce specific goals related with each one of the domains that are to be used in the model. Steps 3, 4, 5, 6, and 7 are part of a design iteration and are intended to produce an increment to the proposed maturity model. Step 3 is focused on designing an architecture for the maturity model. In step 4, we design the increment by selecting which domain will be the focus – Security, AI, DevOps or Governance and its capabilities. In step 5, the focus will be on demonstrating the usefulness of the model for achieving improvement in the target organizations.

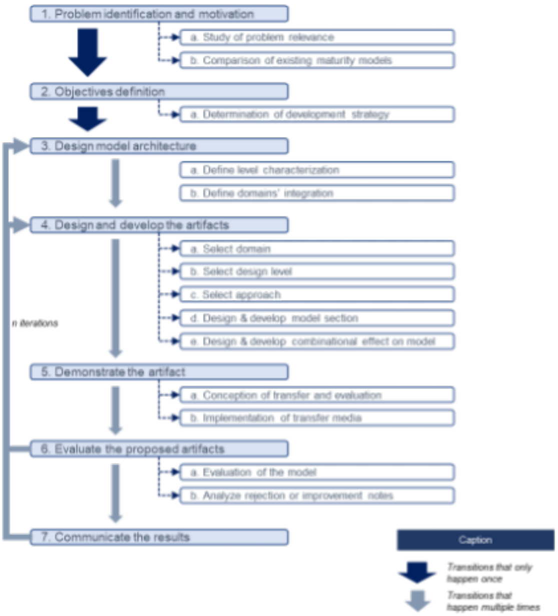
1. Problem identification and motivation
a. Study of problem relevance
b. Comparison of existing maturity models
2. Objectives definition
a. Determination of development strategy
3. Design model architecture
a. Define level characterization
b. Define domains' integration
4. Design and develop the artifacts
a. Select domain
b. Select design level
c. Select approach
d. Design & develop model section
e. Design & develop combinational effect on model
n iterations
5. Demonstrate the artifact
a. Conception of transfer and evaluation
b. Implementation of transfer media
6. Evaluate the proposed artifacts
a. Evaluation of the model
b. Analyze rejection or improvement notes
7. Communicate the results
Caption
Transitions that only happen once
Transitions that happen multiple times

Fig. 1. - Methodology for building the proposed maturity model

For step 6, we have the maturity model evaluated by the industrial partners, gathering quantitative and qualitative data. From the collected data we will produce an analysis on possible improvements and new practices that will be taken into consideration for the next maturity model increment iteration. As a result, we might need to evaluate the model architecture, therefore returning to step 3 or to re-evaluate the design decisions, returning to step 4. In step 7 we communicate the results obtained for the executed increment and the accumulated model produced until then. In this step, we follow an approach similar to the Agile Releases in Software Engineering, where at the end of each incremental cycle the team decides if the product is ready for releasing to end users. Each increment is “releasable” but the decision to release it depends on the value that the overall set of increments can deliver to the users. Likewise, in terms of communication, we will evaluate at the end of each increment if the obtained results are worth communicating *per se* or if they need another increment to be meaningful for the scientific community.

4 Past work and preliminary results

At this point we have already executed steps 1, 2 and 3 from the methodology, which resulted in the presented SLR and on a definition of the model architecture. In this architecture, each domain is characterized by a set of capabilities distributed across maturity levels. Each domain also has an open number of levels, aligned with the objective of producing an open-architecture model. Additionally, to the individual domains’ capabilities and practices, we also need to consider how the different domains interact and influence each other and as such, three types of relationships are defined:

- Type A – Capability to Capability: a relationship between capabilities means that one single capability will immediately activate the need for another, therefore constituting one restriction to the model adaptation.
- Type B – Level to Level: a relationship between practices of different capabilities in the same maturity level, which represents the connection from one practice to another in one same level.
- Type C – Foundational dependencies: a relationship between practices of different maturity levels, which represent a dependency for a higher level.

When applying the maturity model to a specific organization, the defined architecture originates a maturity profile, specific to the organization’s context. To produce it, three steps will be taken. First, the organization gathers data to characterize its initial state and identifies the needs for each one of the domains in the model. Second, the model is applied, making sure that all the domains, dependencies, and restrictions are applied. Finally, as a result, an organizational maturity profile is produced, detailing a set of domains, capabilities, and levels that are to be applicable to that specific organizational context. Additionally, we are currently executing the first iteration for producing the model, which comprises steps 4, 5 and 6 represented in **Fig. 1**.

5 Future work and expected results

The first iteration of our methodology is underway. This work will result in a first definition of the proposed model, which addresses Cybersecurity over DevOps. It will be validated by two different audiences. The first, composed of subject matter experts from academia, with expertise in software engineering and process management. The second, belonging to an industrial partner, a Portuguese telecommunications company, currently evolving its product portfolio with AI infusion and an increased cybersecurity. Once the first validation is completed, we will apply the model in this industrial partner to evaluate its usefulness and applicability in such context.

The following steps would be to execute another iteration and therefore integrate both AI and Governance on the cumulative results.

As we move forward, we expect to achieve a model that is aligned with the defined objectives. This means having a model that can be useful for organizations producing AI-infused software while addressing cybersecurity challenges. Furthermore, having a model with an architecture that can accommodate new domains, as for example, User Experience (UX) or People Management. With this open structure, this model can address the critiques being made to maturity models and therefore, be seen as an improvement tool that different organizations can use.

6 Conclusions

Given the growing importance of software in every aspect of the current societies, even being part of critical infrastructures, software became not just a strategic asset for businesses, but also an attack target for someone trying to get advantage over the others.

With the current work, we produced the first version of a maturity model that addresses known limitations of existing maturity models. First, this proposed model can be used in organizations addressing cybersecurity challenges and producing AI-infused software, with an open-architecture and therefore, able to integrate those two different domains. Furthermore, it is a context-aware model, able to be used in different organizations, given its specific goals. These two characteristics – open-architecture and context-aware are important as it allows for this model to be considered as a useful improvement tool that can be used by different organizations but that can also be evolved by them, addressing the main critique made to the current maturity models as being static and closed.

In this first iteration of our methodology, we focused on integrating cybersecurity capabilities from existing models to make it context aware. In the next iterations, we will focus our work on integrating capabilities and practices needed for producing AI-infused software. Since our SLR showed the inexistence of maturity models in this domain, we will need to identify and characterize new capabilities and practices. Because this is a domain where new knowledge is produced very frequently, this will be a challenge.

While the produced first iteration is already scheduled to be evaluated by one industrial partner, we are also identifying other additional partners that could help us validate the model and its usefulness for different organizational contexts. With these additional partners, we will gather quantitative and qualitative data that will give us the necessary validation for our proposed model.

For future work, once this model can address AI-infused software and cybersecurity, we already envision two topics that can be integrated into the model. The first is the use of UX capabilities and practices, and the second, the use of People Management capabilities and practices. The open architecture for the proposed maturity model will be tested by integrating these new domains without any architectural change.

7 References

1. Marc Andreessen (2011) Why Software Is Eating the World. In: Andreessen Horowitz. https://a16z.com/2011/08/20/why-software-is-eating-the-world/. Accessed 13 May 2022
2. Ransbotham S, Khodabandeh S, Fehling R, LaFountain B, Kiron D (2019) Winning with AI. MIT Sloan Manag Rev 61180:
3. Frijns P, Bierwolf R, Zijderhand T (2018) Reframing Security in Contemporary Software Development Life Cycle. 2018 IEEE International Conference on Technology Management, Operations and Decisions (ICTMOD), Technology Management, Operations and Decisions (ICTMOD), 2018 IEEE International Conference on 230–236
4. Forsgren N, Humble J, Kim G (2018) Accelerate: The Science of Lean Software and DevOps: Building and Scaling High Performing Technology Organizations. IT Revolution Press
5. O'Reilly B (2019) Why Maturity Models Don't Work. https://barryoreilly.com/explore/blog/why-maturity-models-dont-work/. Accessed 12 Jul 2022
6. Dustin Smith, Daniella Villalba, Michelle Irvine, Dave Stanke, Nathen Harvey (2021) State of DevOps 2021
7. da Silva LM, Souza Neto J (2014) Method for Measuring the Alignment Between Information Technology Strategic Planning and Actions of Information Technology Governance. Journal of Information Systems and Technology Management 11:131–152. https://doi.org/10.4301/S1807-17752014000100008
8. ITIL | IT Service Management | Axelos. https://www.axelos.com/certifications/itil-service-management/. Accessed 2 Sep 2022
9. COBIT | Control Objectives for Information Technologies | ISACA. https://www.isaca.org/resources/cobit. Accessed 2 Sep 2022
10. CMMI Institute CMMI. https://cmmiinstitute.com/. Accessed 13 May 2022
11. Wagire AA, Joshi R, Rathore APS, Jain R (2021) Development of maturity model for assessing the implementation of Industry 4.0: learning from theory and practice. Production Planning & Control 32:603–622
12. Blank W, Weitzel J, Blau G, Green SG (1988) A Measure of Psychological Maturity. Group Organ Manag 13:225–238. https://doi.org/10.1177/105960118801300208
13. Becker J, Knackstedt R, Pöppelbuß J (2009) Developing Maturity Models for IT Management. Business & Information Systems Engineering 1:213–222. https://doi.org/10.1007/s12599-009-0044-5
14. Garousi V, Felderer M, Mäntylä M v. (2019) Guidelines for including grey literature and conducting multivocal literature reviews in software engineering. Inf Softw Technol 106:101–121. https://doi.org/10.1016/J.INFSOF.2018.09.006
15. Kitchenham B (2004) Procedures for Performing Systematic Reviews. Keele

16. Adams RJ, Smart P, Huff AS (2017) Shades of Grey: Guidelines for Working with the Grey Literature in Systematic Reviews for Management and Organizational Studies. International Journal of Management Reviews 19:432–454. https://doi.org/10.1111/ijmr.12102

17. Le NT, Hoang DB (2017) Can maturity models support cyber security? 2016 IEEE 35th International Performance Computing and Communications Conference, IPCCC 2016. https://doi.org/10.1109/PCCC.2016.7820663

18. de Bruin R, von Solms SH (2016) Cybersecurity Governance: How can we measure it? In: 2016 IST-Africa Week Conference. IEEE, pp 1–9

19. Carcary M, Doherty E, Conway G (2019) A Capability Approach to Managing Organisational Information Security. Proceedings of the European Conference on Cyber Warfare & Security 97–105

20. Kour R, Karim R, Thaduri A (2020) Cybersecurity for railways – A maturity model. Proc Inst Mech Eng F J Rail Rapid Transit 234:1129–1148. https://doi.org/10.1177/0954409719881849

21. US Department of Energy (2022) Cybersecurity Capability Maturity Model (C2M2)

22. Rodney Petersen, Danielle Santos, Matthew C. Smith, Karen A. Wetzel, Greg Witte (2020) Workforce Framework for Cybersecurity (NICE Framework)

23. National Institute of Standards and Technology (2018) Cybersecurity Framework | NIST. https://doi.org/10.6028/NIST.CSWP.04162018

24. Department of Defense Securing the Defense Industrial Base - CMMC 2.0. https://www.acq.osd.mil/cmmc/index.html. Accessed 7 Feb 2022

25. BSIMM Building Security In Maturity Model (BSIMM). https://www.bsimm.com. Accessed 7 Feb 2022

26. OWASP Project Software Assurance Maturity Model. https://www.opensamm.org/. Accessed 15 May 2022

27. Alsheibani S, Cheung Y, Messom C (2019) Towards An Artificial Intelligence Maturity Model: From Science Fiction To Business Facts. PACIS 46

28. Desouza K, Götz F, Dawson GS (2021) Maturity Model for Cognitive Computing Systems in the Public Sector. Proceedings of the 54th Hawaii International Conference on System Sciences 2173

29. Akkiraju R, Sinha V, Xu A, Mahmud J, Gundecha P, Liu Z, Liu X, Schumacher J (2020) Characterizing Machine Learning Processes: A Maturity Framework. In: Proceedings for Business Process Management: 18th International Conference. pp 17–31

30. Candelon F, di Carlo R, de Bondt M, Evgeniou T (2021) AI Regulation Is Coming. Harv Bus Rev

31. (2021) Europe fit for the Digital Age: Artificial Intelligence. European Commission - European Commission

32. Mora-Cantallops M, Sánchez-Alonso S, García-Barriocanal E, Sicilia M-A (2021) Traceability for Trustworthy AI: A Review of Models and Tools. Big Data and Cognitive Computing 5:20. https://doi.org/10.3390/bdcc5020020

33. Dang Y, Lin Q, Huang P (2019) AIOps: Real-World Challenges and Research Innovations. In: 2019 IEEE/ACM 41st International Conference on Software Engineering: Companion Proceedings (ICSE-Companion). pp 4–5

34. Peffers K, Tuunanen T, Rothenberger M, Chatterjee S (2007) A design science research methodology for information systems research. Journal of Management Information Systems 24:45–77

35. Hevner AR, March ST, Park J, Ram S (2004) Design Science in Information Systems Research. MIS Quarterly 28:75–105