Mine Action Data Observatory: Tackling Uncertainty and System Longevity with Evolutionary Architecture

Juraj Dončević, Kristina Ćaleta, Marko Horvat, Ivica Botički
Faculty of Electrical Engineering and Computing, University of Zagreb
Department of Applied Computing
Unska 3, 10000 Zagreb
`{juraj.doncevic, kristina.caleta, marko.horvat3, ivica.boticki}@fer.hr`

Abstract. *This paper demonstrates the applicability of the concept of evolutionary architectures in the design process for a Mine Action Data Portal. Mine action refers to a range of activities aimed at reducing the threat posed by landmines, unexploded ordnance (UXO), and other explosive remnants of war (ERW). These activities typically include mine clearance, and raising awareness of the general public. The primary design concern for such a system is for it to be long-lasting. The design approach aims to future-proof the system under study by determining possible changes that could affect the system in the future. This is tackled by determining exact design choices as responses to the inferred changes. To avoid the danger of analysis-paralysis during the design process, the design approach acknowledges the existence of uncertainties and circumvents them by introducing adaptivity to the architectural design.*

Keywords. software architecture, evolutionary architecture, data platform, software design, mine action

1 Introduction

Mine action (MA) represents a multifaceted endeavor intertwining humanitarian assistance and developmental studies, focusing on eradicating landmines and mitigating the socio-economic and environmental repercussions of landmines and explosive remnants of war (ERW) (Maathuis, 2003). The overarching objective of mine action is to identify and mitigate the effects and risks posed by explosive hazards to ensure the safety of individuals. Beyond mere demining operations, mine action encompasses a spectrum of measures designed to safeguard communities, aid victims in achieving self-sufficiency and integration, and foster conditions conducive to stability and sustainable progress. Humanitarian demining, inclusive of activities such as mine and ERW surveys, land release procedures, mapping, marking, and clearance, constitutes a vital subset within the broader ambit of mine action endeavors (Bajić et al., 2011) (Matic et al., 2014).

A Mine Action Data Observatory is envisioned as a data portal supporting public information access, MA at operational and strategic levels, and informed decision-making. Its primary task is to manage and analyze vast amounts of data collected from mine-affected regions. The data portal could serve as a critical tool in not only storing and organizing data but also in processing it through advanced methods like optical character recognition (OCR), handwriting recognition (HWR), and information discovery-oriented artificial intelligence (AI) to extract actionable insights that guide mine clearance operations and policy-making. Previous research has already investigated the extent to which modern technological solutions can be used to support MA. It has been determined that ontologies (Horvat et al., 2024) (Horvat et al., 2022) and data lakes (Horvat et al., 2023b) are workable warehousing solutions for MA data observatories. Introducing blockchain technology to MA data observatories has also been proposed to offer tamper-proof data, trust, and transparent data access (Horvat et al., 2023a).

The Mine Action Data Observatory addresses the critical need for integrating diverse and heterogeneous data sources to enhance the efficiency and accuracy of humanitarian demining efforts. In mine action, the complexity and diversity of data gathered through different remote sensing methods, such as multispectral and hyperspectral imagery, ground-penetrating radar data (GPR), magnetic sensor readings, and unstructured documents like hand-drawn minefield maps and accident reports, present significant challenges in terms of data standardization and integration (Krtalić and Bajić, 2019)(Ibrahim et al., 2021). By utilizing formal knowledge representation methods such as the MINEONT+ ontology (Horvat et al., 2022), the proposed observatory semantically integrates these diverse data sets, transforming them into a structured and interoperable format. This approach simplifies the process of data consolidation and enables more efficient retrieval and analysis of relevant information, resulting in more efficient and streamlined decision-making in demining operations in terms of cognitive load and overall operating effort (Bajic, 2010) (Meurer et al., 2010) (Krtalić et al., 2018). To ensure the longevity and adaptability of the Mine Action Data Observatory, an evolutionary architecture is employed, emphasizing flexibility, scalability, and resilience, allowing the system to adapt to new technologies and changing requirements over time. This paper explores how evolutionary architecture can address these challenges, enhancing the observatory's adaptability, scalability, and resilience, thus ensuring its effectiveness and longevity in the dynamic field of humanitarian demining (Horvat et al., 2024).

The principles of evolutionary architecture (Ford et al., 2017) form the basis for the approach; allowing the isolation of changes to specific areas of the system without affecting others, facilitating smoother updates and enhancements.

The remainder of the paper is organized as follows; Section 2 offers the scope and context of an MA data observatory that will be the system under study in the design process. The Observatory's architecture is designed by focusing on system evolvability in Section 3. Future vision and future work are discussed in Section 4. Section 5 provides the conclusion to the paper's proposal.

2 Observatory scope and context

The Observatory's domain is centered around mine suspected areas (MSA) and the assortment of data gathered about them through mine records, interviews, monographs, MAs, MSA indicators, MSA interpretations by experts, and monographs (Fig.1). The presented entities are laden with geographical data, containing precise positioning of mission-critical objects and areas. The domain doesn't contain any structural complexities or elaborate branchings from the central MSA entity, therefore no detachments are envisioned in its physical implementation.

The proposed Observatory is envisioned as a data portal to facilitate both past and future data. The use cases for the Observatory outline two general purposes for the system; as an archival system handling historical data and an operational system handling recent data. The data is expected to be ingested into the observatory through a data portal which is implemented to support heterogeneous data ingestion.

The context of the system is presented in Fig 2., according to the C4 design model (Brown, 2010) (Richards and Ford, 2020). CMAC is expected to enter existing MA data into the observatory. This data is expected to predominantly consist of scanned physical documents of hand-filled forms such as mine action reports, interview minutes, testimonies, monographic documents, and GIS digital files. The greatest challenge in supporting this ingress data flow is combining object detection, optical character recognition (OCR), and handwriting recognition (HWR) to enable semi-autonomous document entry. These methods are required to ease the archival document entry process for system operators.

Data is also expected to ingress the observatory as current operational MA reports are entered by demining companies. The ingressing data should be entered through forms, enforcing standardization of both the data structure and format.

Consumed ingress data must be standardized regardless of its archival or operational nature - it is distilled into a unified model to support multiple consumption cases. The standardized data is expected to be used by MSA interpreters possessing expert domain knowledge. Interpreters are tasked with enriching the existing data and providing further insight into MA data. This includes expanding or correcting MSAs, or providing additional semantics to raw data to guide future MAs.

The general public will be served MA data through a read-only data portal that supports GIS visualization. Due to the expected sensitivity of some data, it should be anonymized and curated. If required, documents should be redacted by blackouts and the accompanying data anonymized.

To facilitate future information and knowledge discovery, the consumed and standardized observatory data will be used to train various AI models. These could be used in future decision support systems for minefield navigation or for guidance in MSA discovery where data is sparse or incomplete.

3 System capabilities and architectural design

Richards and Ford (Richards and Ford, 2020) claim that prioritizing system architecture is an effective way to tackle some requirements placed on a piece of software. Requirements are allowed to freely project themselves onto the architectural design of the Observatory.

The key requirement impacting the Observatory's architectural design is the need for the system to be long-lived due to its archival nature. This does not just encompass the deployability of the executable code, but also the maintainability of the Observatory's codebase. The primary goal is to implement a system that can be kept updateable, modifiable, and deployable not just in terms of years, but of decades. The Observatory must be able to evolve during its lifespan by tackling technological and domain changes; marked with the architectural meta-characteristic of evolvability. Evolvability is achieved by aligning a system architecture with the concept of evolutionary architecture (Ford et al., 2017).

Our research has determined that the dimensions of change that might impact the Observatory are:

- Changes in web or client-side application technology stacks;
- Changes in framework technology;
- Changes in data storage technologies;

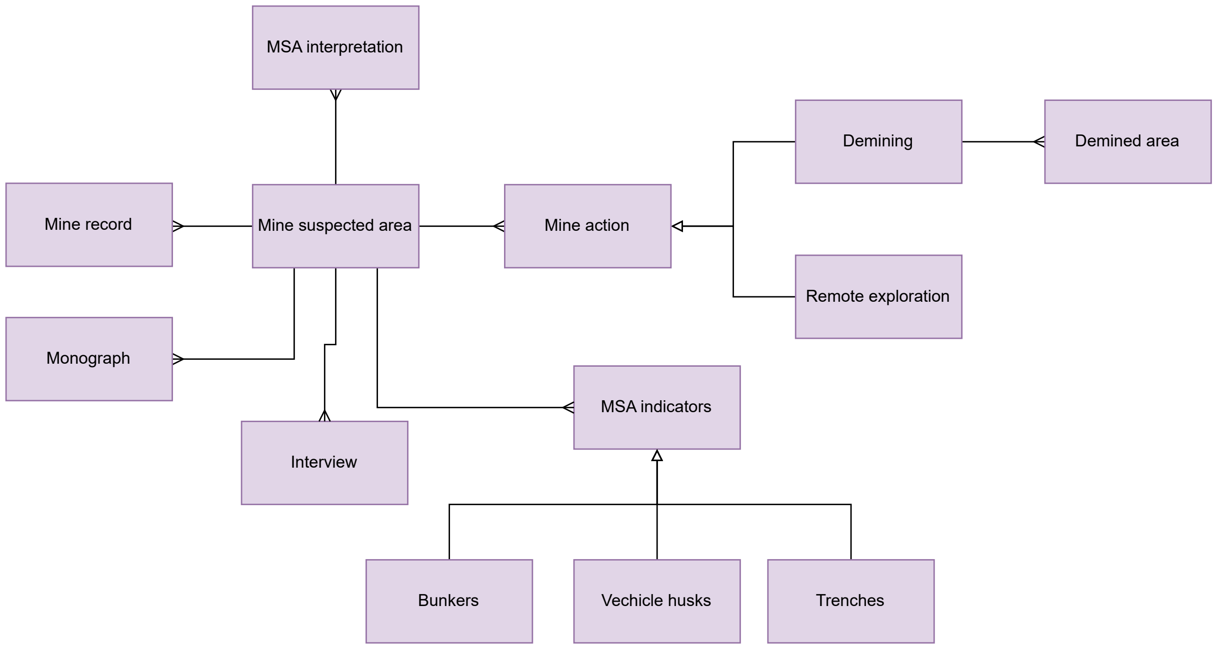

Figure 1: A simplified conceptual domain model for MA

- Changes in the document formats that need to be processed by the system;
- Changes in AI technology stacks and tools;
- Changes in data warehousing paradigms;
- Changes in load switching between operational and archival system context parts.

Determining impactful changes has provided a unique opportunity to push ahead the system architecture design in the overall design process. Along with the general requirements, the only specific detail known at the time was the conceptual domain model and that it didn't require the separation of the domain into separate bounded contexts - necessitating the decoupling within the domain model itself (e.g. via vertical slicing or microservices).

The architectural design was conducted earlier in the project life-cycle than usual but with an initiative to accommodate any changes in requirements without constraining the project. This allows managers, researchers, and developers to determine both certainties and uncertainties regarding the shape of the future system.

The certainties were represented by the information that the system would contain a separately implementable and deployable core service, a database, a file storage solution, and a previously implemented ontology knowledge database.

The uncertainties were tackled by separating the operational concerns of the Observatory into individual components. A general decision was made to treat the unspecified components as implementation details. The impact and propagation of their changes throughout the system would be minimized by interfacing. The interfaces are not determined ahead of the implementation process but are inferred only when implementing a specific functional point requires their actualization. Some notable uncertainties faced during the design were the technology for the client application, the technology of the file storage solution, and the specific technological implementation and time requirements of the HWR/OCR and AI model services.

The following design decisions were made to tackle the aforementioned changes that the system might have to evolve around:

- The client-side application was extracted into a separate executable component to support changes without impacting the core service.
- A core library and domain were implemented without 3rd party dependencies; effectively decoupling the domain from the framework used to implement the core service. The core domain is propped up by concepts of domain-driven design (DDD) (Evans, 2003) and functional programming (FP). DDD principles were used to design the domain and to characterize models as entities, aggregates or value objects. DDD effectively provided a systematization of the design process and detached domain considerations from technological specifics. FP was used to enable monadic chain calls, the control of exceptions and errors and the expression of nullability. In this way, technological features were detached from the implementation of core workflows, data exchange, framework usage, and error handling.
- All technological changes (including changes in data storage technology) are supported by the ports-and-adapters pattern (Cockburn, 2005). An interface rep-

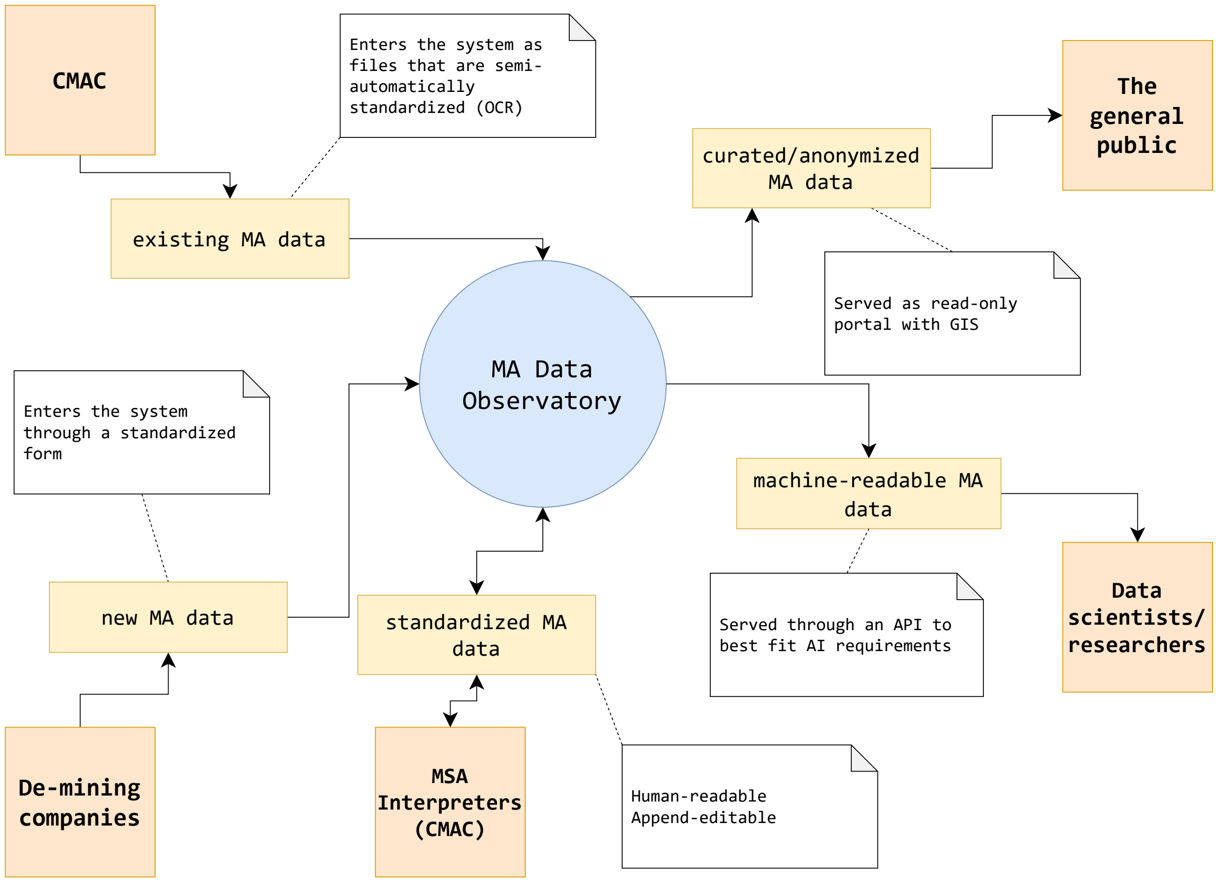

Figure 2: Observatory system context and data flow

resents a port, while its implementations represent adapters that can be swapped out by dependency injection.

- Changes in document processing are primarily supported by the extraction of OCR and HWR services from the core service; allowing the use of strategies for specialized physical document analysis.

- All AI-oriented solutions are extracted from the core service, allowing them to be deployed on specialized hardware optimized for machine learning (e.g. FPGA boards, GPUs, specialized CPUs).

- Data warehousing functionalities are extracted from the core service to allow technological flexibility of the warehousing solution. An unstructured data lake in formats like Avro or a structured star schema warehouse is expected.

- Since it was impossible to determine the computational load required from the core service at design time, the core service was split into two vertical slices (Newman, 2019). The split was not influenced by the domain, but by the dual system use cases - archival or operational system. Due to the possibility of the slices being split into separate microservice components, the WET (write everything twice) principle (Pai and Xavier, 2017) was adopted for everything except the core library.

Although evolutionary architecture proposes the use of event-driven architecture, this was not followed in this case to simplify orchestration due to the relatively small deployment scale (e.g., no need for service replication or redundancies). Conversely, an orchestration or choreography pattern would introduce architectural complexity in terms of an additional mediator component for the former, or multiple message queues in the latter case. Due to the context of the system, we found this trade-off to be unjustifiable, as the added complexity would not provide a significant benefit over the current design. Despite this, message queues are still utilized to provide asynchronicity for time-exhaustive operations. The case-specific message queues are also introduced to the core service through the ports-and-adapters pattern.

Fig 3. represents the container part of the C4 model and illustrates the described architectural design decisions. Although the system may be adapted to the cloud due to its containerization, its initial deployment might not be on the cloud. A simpler goal has been set for a proof of concept - the ability to deploy the system on a single via a Docker compose script. The system design even future-proofs the containerization technology by providing the ability to configure the system

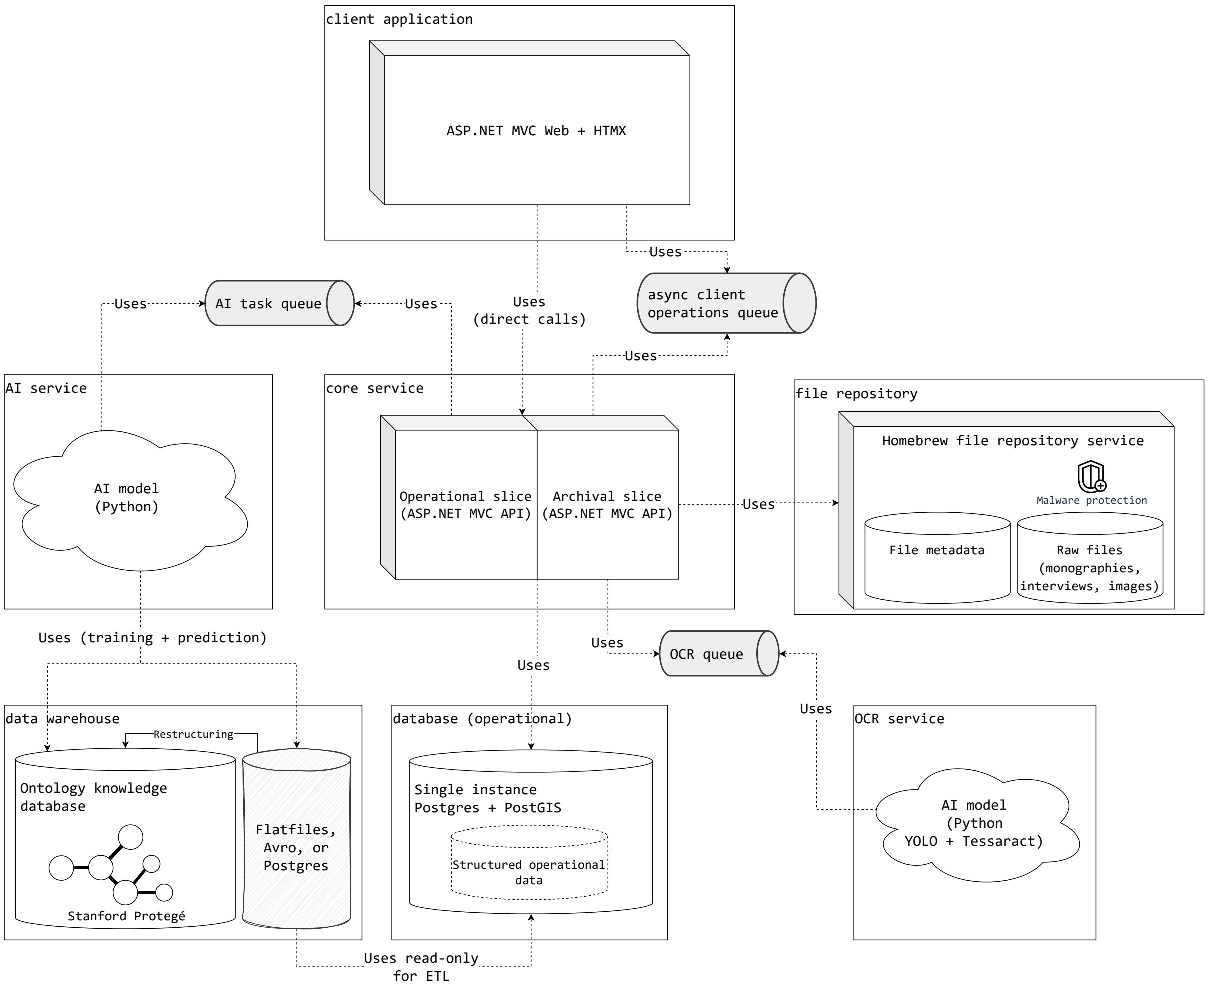

Figure 3: Observatory containers and component outline (and proposed technologies)

components through configuration files.

4 Future vision and future work

The Observatory system is designed to be robust, adaptable, and responsive to continuous change through evolvability. Future work will focus on the continuous observation of new technologies and trends in frameworks, data storage, AI, and OCR/HWR technologies. The Observatory offers a playground to test emerging technologies, as well as the ability to test its own evolvability. The assessment of these results through future research will provide more insight into appropriate strategies and techniques for future-proofing software systems.

It is expected that the Observatory, if successfully implemented, will not only meet the archival challenges in Croatia, but also in other countries with unsupervised MSAs or unfortunate countries with ongoing conflicts. The latter will be the real test of the Observatory's operational viability and longevity.

5 Conclusion

The development of the Mine Action Data Observatory represents a significant achievement in the realm of mine action initiatives, marking a decisive step towards a safer and more sustainable future. Through strategic design decisions following evolutionary architecture principles, a system can be crafted that not only addresses the immediate challenges of mine action but also embodies resilience and adaptability in the face of uncertainty. By designing a long-term solution resistant to change, we have fortified the observatory against unpredictable shifts in operational landscapes, technological advancements, and user requirements, ensuring its enduring relevance and effectiveness.

Furthermore, the Observatory's comprehensive scope underscores its pivotal role in addressing the multifaceted dimensions of mine action data. By standardizing data about demining operations the observatory empowers stakeholders with invaluable insights and resources to make informed decisions. Its capacity to facilitate transparent and accountable mine action efforts sets a new standard for the field, promoting best practices and driving positive outcomes for affected communities worldwide.

The Observatory stands as a testament to the power of technology in support of lasting humanitarian endeavors. As challenges and opportunities emerge, the Observatory will continue to be incrementally designed and implemented. Work focusing on refining the system design in regard to deployment capabilities and orchestration is expected as it gets deployed in different varieties and physical locations. The Observatory will also provide a playground for experimenting with state-of-the-art technologies and principles.

Acknowledgments

This work was supported by the Ministry of Science and Education of the Republic of Croatia grant: NPOO.C3.2.R3-I1.01.0268.

References

Bajic, M. (2010). The Advanced Intelligence Decision Support System for the Assessment of Mine-suspected Areas. *The Journal of Conventional Weapons Destruction*, 14(3).

Bajić, M., Matić, , Krtalić, A., and Vuletić, D. (2011). Research of the Mine Suspected Area. *Research of the Mine Suspected Area*.

Brown, S. (2010). The C4 model for visualising software architecture.

Cockburn, A. (2005). Hexagonal architecture.

Evans, E. (2003). *Domain-Driven Design: Tackling Complexity in the Heart of Software*. Addison-Wesley Professional, Boston, 1st edition edition.

Ford, N., Parsons, R., and Kua, P. (2017). *Building Evolutionary Architectures: Support Constant Change*. O'Reilly Media, Beijing, 1st edition edition.

Horvat, M., Krmpotić, K., Krtalić, A., and Akagić, A. (2023a). Bridging Blockchain Technology and Humanitarian Demining: A Novel Concept for Decentralized Storage of Landmine and UXO Locations. pages 369–375. ISSN: 1847-2001.

Horvat, M., Krtalić, A., Akagić, A., Krmpotić, K., and Skender, S. (2023b). Humanitarian Demining Using Data Observatories and Data Lakes. *19th International Symposium "Mine Action 2023", Book of Papers*, page 31.

Horvat, M., Krtalić, A., Akagić, A., and Mekterović, I. (2024). Ontology-Based Data Observatory for Formal Knowledge Representation of UXO Using Advanced Semantic Web Technologies. *Electronics*, 13(5):814. Number: 5 Publisher: Multidisciplinary Digital Publishing Institute.

Horvat, M., Krtalić, A., Bajić, M., Muštra, M., Laura, D., and Gold, H. (2022). MINEONT: A proposal for a core ontology in the aerial non-technical survey domain. pages 47–51. ISSN: 1849-3718.

Ibrahim, N., Fahs, S., and Alzoubi, A. (2021). *Land cover analysis using satellite imagery for humanitarian mine action and ERW survey*.

Krtalić, A. and Bajić, M. (2019). Development of the TIRAMISU Advanced Intelligence Decision Support System. *European Journal of Remote Sensing*, 52:40–55.

Krtalić, A., Racetin, I., and Gajski, D. (2018). The Indicators of Mine Presence and Absence in Airborne and Satellite Non-technical Survey. pages 57–61.

Maathuis, B. H. (2003). Remote Sensing Based Detection of Minefields. *Geocarto International*, 18(1):51–60. Publisher: Taylor & Francis _eprint: https://doi.org/10.1080/10106040308542263.

Matic, , Laura, D., Turšić, R., and Krtalić, A. (2014). *Analytical assessment for the process of collecting additional data on a suspected hazardous area in humanitarian demining*.

Meurer, H., Wehner, M., Schillberg, S., Raven, N., Wirtz, T., and Ime, F. (2010). An Emerging Remote Sensing Technology and its Potential Impact on Mine Action.

Newman, S. (2019). *Monolith to Microservices: Evolutionary Patterns to Transform Your Monolith*. O'Reilly Media, Beijing China ; Sebastopol, CA, 1st edition edition.

Pai, P. and Xavier, S. (2017). *.NET Design Patterns*. Packt Publishing Ltd. Google-Books-ID: t08oDwAAQBAJ.

Richards, M. and Ford, N. (2020). *Fundamentals of Software Architecture: An Engineering Approach*. O'Reilly Media, Sebastopol, CA, 1st edition edition.