

Model Driven Development for AI-based Healthcare Systems: A Review

Colm Brandon^{1,2}[0000-0001-7842-7897]
Amandeep Singh^{1,2}[0000-0002-6371-0435]
Tiziana Margaria^{1,2,3,4}[0000-0002-5547-9739]

¹ University of Limerick, Limerick, Ireland

² Centre for Research Training in Artificial Intelligence (CRT-AI), Ireland

³ Lero, Ireland

⁴ The Health Research Institute (HRI), University of Limerick, Ireland

Abstract. We review our experience with integrating Artificial Intelligence (AI) into healthcare systems following the Model-Driven Development (MDD) approach. At a time when AI has the potential to instigate a paradigm shift in the health sector, better integrating healthcare experts in the development of these technologies is of paramount importance. We see MDD as a useful way to better embed non-technical stakeholders in the development process. The main goal of this review is to reflect on our experiences to date with MDD and AI in the context of developing healthcare systems. Four case studies that fall within that scope but have different profiles are introduced and summarised: the MyMM application for Multiple Myeloma diagnosis; CNN-HAR, that studies the ability to do AI on the edge for IoT-supported human activity recognition; the HIPPP web based portal for patient information in public health; and Cinco de Bio, a new model driven platform used for the first time to support a better cell-level understanding of diseases. Based on the aforementioned case studies we discuss the characteristics, the challenges faced and the positive outcomes achieved.

Keywords: Model Driven Development · Health Informatics · Artificial Intelligence Applications · Low-Code/No-Code · Digital Thread

1 Introduction

The integration of AI into healthcare systems is potentially a paradigm shift for the healthcare sector, with impacts on a variety of levels, such as computer aided diagnostics, personalised medicine, drug discovery, disease understanding, and the processing of healthcare records to name a few. AI's ability to extract actionable knowledge from large amounts of unstructured data has huge opportunities to advance medical discovery and also free medical professionals from a variety of administrative tasks through automation.

There are also potential pitfalls. In computer aided diagnostics, implicit or explicit biases in the training data could lead to higher rates of misdiagnosis

in underrepresented patient groups. There is a trade-off between preserving patient privacy and maximising treatment outcomes with respect to personalised medicine. It is therefore critical to embed medical professions and other stakeholders in the development process as much as possible.

MDD can be a way to include professionals more and earlier in the development of new healthcare software and systems. MDD works using models of a given domain, be they models of organs (like a digital twin), of procedures (like workflows) or of knowledge (like databases, ontologies, and rules). Some of these models are represented as data, others as graphs (like workflows and processes) which can be later transformed into executable code, human or AI generated. These models capture the essence of the domain and the problem the system solves. Most of them enable the creation of Domain-Specific Languages (DSLs) which express the functionalities in the domain in a more intuitive way for the health professionals, and this way improve the usability and effectiveness of the models and the systems. The key property is that models and DSLs must offer a level of abstraction directly accessible to the domain experts, shielding the technicalities. The goal is that domain users may be able to understand the system’s design and functioning without the need to also understand the underlying technologies for the implementation and runtime.

This paper explores our experiences to date with respect to the two dimensions of *advanced model driven development* and *AI* in the development of healthcare systems. We summarise four case studies with different characteristics in this spectrum, covering different applications in healthcare: diagnosis in clinical medicine (MyMM), primary care for neurological impairment (CNNHAR), patient information in public health (HIPPP) and the cell-level understanding of diseases (CINCO de Bio (CdB)).

The paper is organized as follows. Section 2 considers literature and previous work pertaining MDD and AI in Healthcare Systems. Section 3 presents and discusses the four cases studies. Section 4 discusses the characteristics, challenges and successes in the cases studies. Finally Section 5 offers some final thoughts on the current work and perspectives for future work.

2 Model Driven Development and AI in Healthcare Systems

We consider here the two main perspectives of AI and MDD in healthcare.

2.1 AI in Healthcare

AI and Machine Learning (ML) are increasingly embedded in healthcare systems all the way from the laboratory, in biomedical research, to clinical medicine and patient care at home.

In health related research, AI models have proven to be extremely effective for the prediction of protein structures [27] in drug discovery, and to predict physiochemical properties [67] of environmental chemicals. When understanding

diseases, AI models enable single cell segmentation in tissues, and thus single cell analysis [23, 37].

In the context of primary care, Internet of Things (IoT) devices and AI at the edge have been used effectively to mitigate the negative side effects (without direct human intervention) for individuals with a variety of illnesses. An excellent example is the use of smart devices to predict and detect gait freezing episode of Parkinsons patients, so that an automated intervention can un-freeze the patient [10, 11].

In public health, AI and ML have been successfully used for information analysis and data mining. During the COVID pandemic, Irish researchers developed a method of estimating the reproduction or ‘R-number’ in COVID-19 cases, the number that describes the likelihood of a virus increasing or decreasing in a population. Their susceptible-exposed-infected-removed – or SEIR – model, based on statistics and ML, was used regularly to provide up-to-date scenario analysis to the National Public Health Emergency Team (NPHET) [21]. Similarly, AI helped identify diseases and foodborne illness outbreaks through data mining social media data [54, 58, 59]. AI has also been used to identify individuals suffering with mental illnesses through the analysis of social media data [30, 48, 62]. AI models have been developed to mitigate the effect of low-quality health information on treatment outcomes of patients, by automatically evaluating the quality of web-based health information with respect to some established evaluation frameworks developed by physicians [31].

In clinical medicine, AI models support clinicians in a variety of ways. Medical image classification models assist in the diagnosis of a variety of diseases [68, 65, 25, 34]. Examples are mammogram classification in the context of breast cancer diagnosis [4], heart disease [46], and many more. The possibility of using Large language Models for computer aided diagnosis is being explored [61], along with the development of clinical decision support systems that use other AI methods [66]. LLMs are also expected to help reduce the ever growing administrative workload.

2.2 Model Driven Development in Healthcare

MDD has been used extensively in the healthcare domain in recent times with impacts across a broad range of applications such as clinical medicine, information systems for healthcare and health monitoring systems.

In the context of clinical medicine, MDD has been used to simplify the process of developing software systems for clinical scoring and mobile health through the use of meta-modelling [18]. A meta-model that facilitates the development of dynamic checklist support systems for clinical safety helps reduce the cognitive load for clinical stakeholders when designing such systems [49].

With respect to information systems for healthcare, a MDD based versatile Nursing Information System (NIS) can cater for a variety of clinical requirements across a range of clinical settings and be deployed on different devices [63]. For general Hospital Information Systems (HIS), a multidimensional supervision metamodel was developed to create a general framework that mitigates

the challenges of developing such systems which have to cater for a wide variety of end-users and healthcare delivery processes [35]. Agile and model driven methodologies have been shown to aid in the development of Electronic Health Record (EHS)-based population registries and their associated data collection tools [28]. Model-Driven Engineering (MDE) principles have also enabled software engineers to develop HL7 (a healthcare information and protocol standard) models using a Unified Modeling Language (UML) based interface [50].

In the context of personal health monitoring, a model-driven computational framework for Personalized Health Monitoring (PHM) applications using IoT devices poses specific emphasis on security and energy issues [60]. A framework for general IoT-based healthcare monitoring systems presents a meta-model which defines the fundamental components of such systems and their relationships [3].

In the specific context of MDD, we use a set of Integrated Modelling Frameworks that are described next.

2.3 Model Driven Development Frameworks

Our approach simplifies the IT and integration tasks by composing *models of behaviour* to achieve the same results. By introducing the different DSLs as the high-level abstraction for each behaviour and plugging them in, we enable a development cycle where stakeholders, instead of only defining requirements, can take action by participating actively in the successive stages of the project. This approach has been successful in several domains, for example in smart agriculture applications [24].

Our contributions cover the demonstration of various aspects of the use of two Low-Code / No-Code (LCNC) development environments, DIME [12] and Pyrus [69], and the extension of their Application Domain Specific Languages (A-DSLs) to cover healthcare settings. To accomplish this, we extend the range of A-DSLs that these two environments provide in order to enable our external data sources and systems to become amenable to our LCNC application development approach, allowing a high level of reusability. This extension involves creating several building blocks to communicate with the different systems, that may as well include IoT devices like the low-cost compact multi-sensor platform Nordic *Thingy:53*⁵ that we use for example in the Human Activity Recognition (HAR) context. The A-DSLs in particular enable analytics pipelines for a subsequent decision-making.

We embrace a LCNC software development paradigm [9], that is rapidly gaining foot in industry and is predicted to become the development style of choice for 80% of newly developed software by 2026 [1]. However, we specifically adopt a Model Driven Design and development paradigm [45, 38] where the models are not just graphically suggestive but also have an underlying formal model in terms of Kripke Transition Systems [33]. This choice makes them analyzable through well-established techniques like control flow and data flow analysis [36], model checking [8], property checking [56], reachability analysis [29] and more,

⁵ https://docs.nordicsemi.com/bundle/ug_thingy53/page/UG/thingy53/intro/frontpage.html

like synthesis [41], [40], also in robotics and IoT contexts [26]. Specifically, we use here the two Integrated Modelling Environments DIME and Pyrus.

DIME is an Eclipse-based LC/NC graphical modelling environment that enables prototyping web applications in a model-driven way. It follows the One Thing Approach (OTA) [43] and the eXtreme Model-Driven Design (XMDD) [44] paradigms for modelling and development, empowering domain experts to model an end-to-end web application with no programming experience. To cover the different aspects of web applications, DIME provides a collection of ready-to-use modelling languages, as well as collections of native DSLs that support the development of new applications via composing models of different natures. In particular, DIME allows users to define the data model, create the user interface models, and create the workflows. Those models are checked for syntactic compliance largely automatically ⁶.

Pyrus is a web application that offers a graphical, collaborative development environment for Data Analytics. It bridges the gap between Python-based established programming platforms like Jupyter [32] and graphical workflow composition in a data-flow fashion. Individual Python functions are implemented and stored in Jupyter, special signature annotations are added to these functions and exported to Pyrus, so that the functions can be identified and retrieved by the Pyrus web-based orchestration tool, where the pipelines are composed. From the pipelines, Pyrus generates the Python code for the orchestration and configuration, which is again stored and executed in Jupyter. This separation of concerns decouples the coding and development of the single functionalities (in Python/Jupyter), the low-code part of the approach, from the data analytics orchestration modelling, which happens in accordance with model-driven engineering principles and completely graphically, in a no-code fashion.

3 Case Studies

Here we summarise four selected case studies that exemplify our experience of bringing model driven development and/or AI into the healthcare domain.

3.1 Automating the Diagnosis of Multiple Myeloma - XMDD and Web

Multiple Myeloma (MM) is a blood cancer that develops in plasma cells in the bone marrow. It is one of the deadliest forms of blood cancer, with a five-year survival rate as low as 40%. In collaboration with the University Hospital Limerick, in the context of the interdisciplinary UL Cancer Network (ULCaN), we automated the referral pathway for MyMM through a Web application developed in DIME, using XMDD.

MM is a cancer with circulating monoclonal paraproteins. It is diagnosed by measuring risk factors in the blood: monoclonal proteins called Paraproteins

⁶ DIME can be downloaded from <https://scce.gitlab.io/dime/>

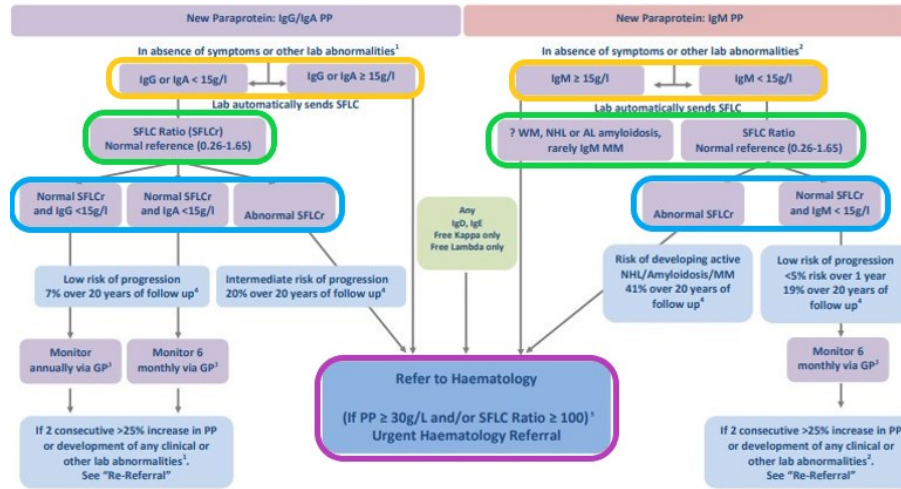


Fig. 1. Extract from UHL Guidance on Management of MGUS in Primary Care [51]: risk stratification workflow.

(PP), and Serum-Free Light chain (SFLC) ratios, in blood samples. When lower levels of paraproteins are detected in healthy individuals, this is known as Monoclonal Gammopathy of Uncertain Significance (MGUS). A risk stratification protocol along the levels of such factors developed at the University Hospital Limerick [52] (see Fig. 1) recommends a patient's further care depending on their outcome. These outcomes can range from a yearly check-up and monitoring via a General Practitioner (GP) to an urgent referral for a patient to be seen by a consultant within the Haematology department. We aimed to improve referral timelines for patients with possible MGUS and automate the risk stratification process by replacing the manual process with a web-based application developed in DIME with XMDD technologies. The MyMM web application implements the manual reference risk stratification algorithm and for patients with possible MGUS it communicates a patient's further care to the requesting physician or GP. The system includes a process model implemented in the DIME development environment.

Fig. 2 shows the corresponding top level decision process modelled in DIME, with a layout resembling the layout of the original algorithm.

It contains several process models that carry out the task of completing each comparison for each layer of the risk stratification seen in figure 1. Care was taken to mimic in the DIME process structure and layout of the original algorithm, which is known to the healthcare specialists from their training.

Solving this particular problem is important because the automation makes the evaluation and communication of outcomes immediate, instead of waiting for days until a trained specialist finds time to manually score the accumulated batch of samples. In case of the worst outcomes, timeliness is crucial.

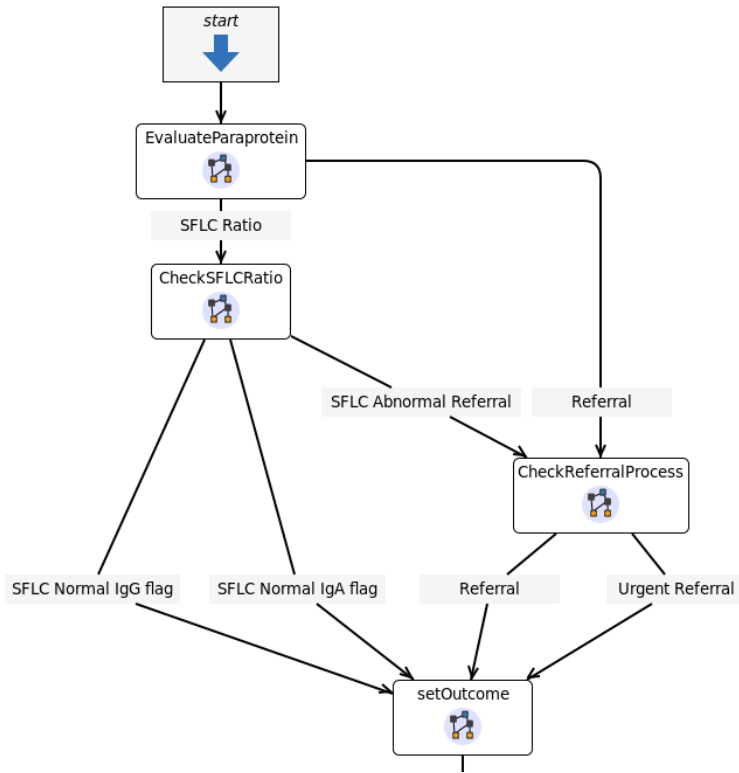


Fig. 2. The top process model of the risk stratification process (in DIME).

A key challenge was the aim to develop generic processes and components, easily parametrizable, that can be easily reused for other risk stratifications or scoring algorithms. The adopted solution is to abstract parameterizable components and develop generic processes.

The full-stack web-application generated and deployed from the DIME models and Service-Independent Building Blocks (SIBs) also facilitates users to interact with the computational risk stratification algorithm through a web-interface. This reduces the need for skilled healthcare staff to manually apply the algorithm to patient data sets. The DIME development environment lends itself to mapping workflows and process control flows, making development in this manner a utility.

3.2 CNN-based Human Activity Recognition on Edge Computing Devices

Research on HAR involves wearable devices integrating inertial and/or physiological sensors to classify human actions and status across various application domains, such as healthcare, sports, industry, and entertainment [19]. However,




Fig. 3. Full monitoring of a HAR subject with wrist sensors and other monitors (from [19])

executing HAR algorithms on remote devices or the cloud can lead to issues such as latency, bandwidth requirements, and energy consumption. Transitioning towards Edge HAR can be a more effective and versatile solution, overcoming the challenges of traditional HAR techniques. A novel HAR model for computation on edge devices was presented in [55], with a Convolutional Neural Network (CNN) Deep Learning approach, then compared with cloud-computing HAR models.

Figure 4 describes the adopted HAR processing pipeline, with all the processing steps: Data collection, Noise removal, Segmentation, Feature extraction, Feature selection, Train/Test split, Training HAR model, Deployment on Edge device, Performance exploitation. The CNN model was trained on the Edge Impulse cloud platform by implementing the pre-processing steps of Figure 4 and deployed on the Thingy 53 device for real-time recognition and evaluation.




Fig. 4. Detailed HAR system overview. (from [55])

The edge computing model achieved promising results ($\geq 92\%$) in terms of Precision, Recall, and F1-score, and with significantly reduced latency and minimal memory. Models with higher computational complexity can be deployed

in the cloud but due to network latency and longer inference times they are not suitable for real-time tasks

As we see, there is here a good outcome for gesture recognition, leveraging CNN-based deep learning, but there is de facto no MDD in the implementation, although the phased pipeline (with sub-workflows) of Fig. 4 would lend itself to the technique.

3.3 HIPPP: Health Information Portal for Patients and Public

This case study [14] concerns the design and implementation of a public-facing web application that facilitates the automated evaluation and "scoring" of web-based health information. Here, MDD is used to create the web application (similar to the MM case), and to create a specific DSL which enables domain users to use AI pipelines to conduct the end-to-end analysis of web based health information with respect to established quality evaluation frameworks [53] developed by physicians. The AI DSL encompasses a variety of AI models across Natural Language Processing (NLP) and graph processing to perform feature extraction and classification, and ultimately to enable the evaluation process to be automated. The HIPPP system was developed as a web application in DIME. HIPPP utilises a variety of AI algorithms (*transformers, graph neural networks, random forests*) and other computational techniques (*dynamic programming*) to implement an end-to-end classification pipeline that automates the application of the QUEST framework (see Figure 5). Due to our collaboration with colleagues in the Health Research Institute who work on colon cancer, it is tailored so far to Web-based Health Information (WBHI) pertaining to the diagnosis, treatment and prevention of colon cancer. The reliability and trustworthiness of the sources are subjected to automated evaluation based on the QUEST framework. In the context of the quest for more and more direct patient inclusion and increased quality of information, the goal is for the system to act as a safeguard against patients and the public unknowingly basing their medical decisions on information that they should not trust.

It is foreseen to have an expert in the loop for cases that are not absolutely clear, and to allow adaptive learning based on the manual scores that the expert(s) provide.

The first outcome of this work is the HIPPP application seen as a ready-made fullstack web application which could be deployed by any public health body or advocacy group and be used by patients and the public to evaluate the information they access, providing it is trained on the specific health conditions of interest. The second outcome is the extension of the capabilities of DIME with an A-DSL for AI-based parsing and scoring of document based information sources. It includes several SIB palettes and several processes (for the pipelines and other algorithms) that are retargetable towards different scoring systems and reusable for other conditions of interest.

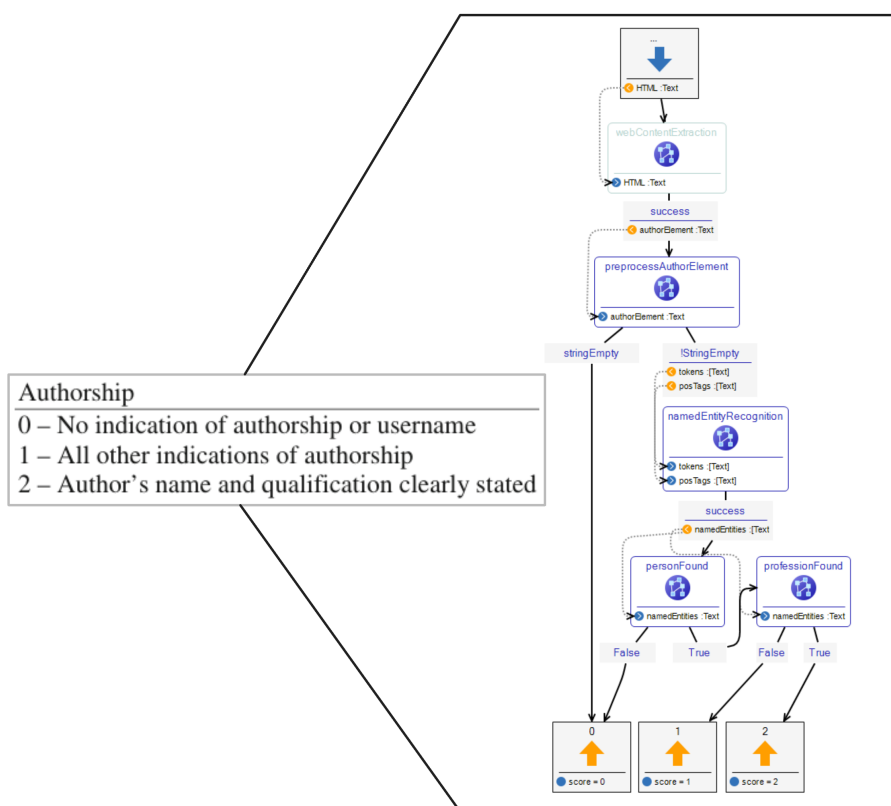


Fig. 5. A condensed view which shows translation of the QUEST authorship criteria classifier (image from [53]) translated to a DIME process model (image from [14]). Illustrating how the authorship scoring can be automated, using a combination of web content extraction, named entity recognition and simple rule-based classifiers.

3.4 Cinco de Bio: A Platform for Domain-Specific Workflows that Leverage AI for Biomedical Research

Cinco de Bio [13] is a platform which enables users to design and execute biomedical analysis workflows using a no-code graphical modelling language. The specific case study concerns cell segmentation carried out on large images of tissue samples, with many layers of information. Fig. 7 summarises the process.

The analysis and segmentation software was already available in an implementation in Python, R and Matlab. We optimised some of the algorithms and refactored it to be amenable to becoming SIBs, and redesigned the workflows, while at the same time developing the Cinco de Bio modelling, compilation, deployment and execution environment.

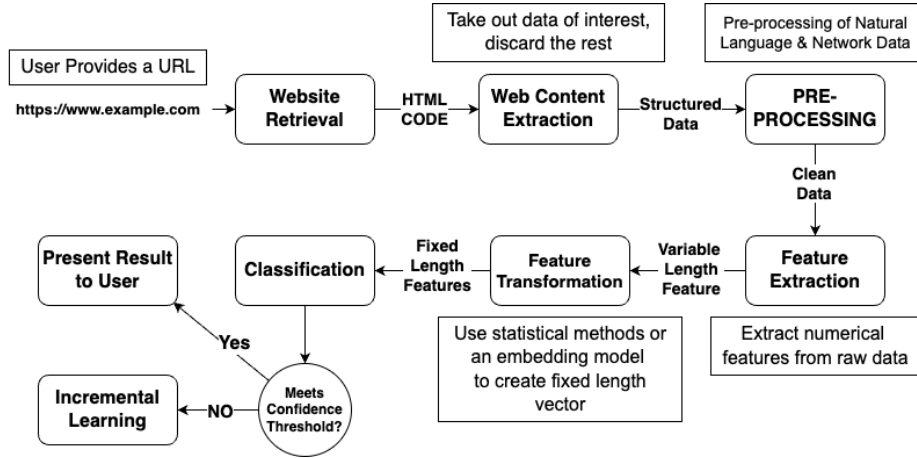


Fig. 6. A high-level overview of the HIPPP classification pipeline (image from [14]). With the developed DIME A-DSL consisting of SIBs which fall under one of the workflow steps namely *web content extract*, *pre-processing*, *feature extraction*, *feature transformation*, *classification* and *incremental learning* for graph and natural language data.

The AI aspect of the project concerns the support in the recognition of single cells, by improving the training capability through artificial image generation. This is currently ongoing work.

The platform comprises a) an Integrated Modeling Environment (IME) for designing and validating workflows built using the Cinco meta modelling tool and b) a workflow execution environment which translates the workflow models to programs that orchestrate the analysis workflows. The execution environment is a Kubernetes-native [15] application that comprises of several custom core services that are loosely coupled. These services handle the execution front-end, model-to-code transformation, computational environment for workflow orchestration program(s), data management, SIB management (which are stored in container-registries (such as Dockerhub)) and a variety of APIs for handling job-scheduling, SIB interactions, etc.. The core services are also accompanied by a number of open source services to deliver the overall platform functionality, namely RabbitMQ [64] (message passing), MinIO [2] (Cloud Native Storage), MongoDB [47] (NoSQL database) and Kaniko (building container images in k8s). In addition to the CdB core platform we have developed a SDK for integrating Python and R-based tools into CdB as SIBs, to ease the tool integration process.

Concerning the A-DSL and the modelling style, the SIB and workflow modelling style in CdB is kept quite similar to the SIBs and processes shown in Fig. 2 for the Myeloma diagnosis case study and Fig. 5 for the HIPPP case study. This is intentional, because the style has proven itself in a wealth of different applications from different domains (smart manufacturing, IoT, AI and reasoning, decision support systems, games, and more), and also because it simplifies sup-

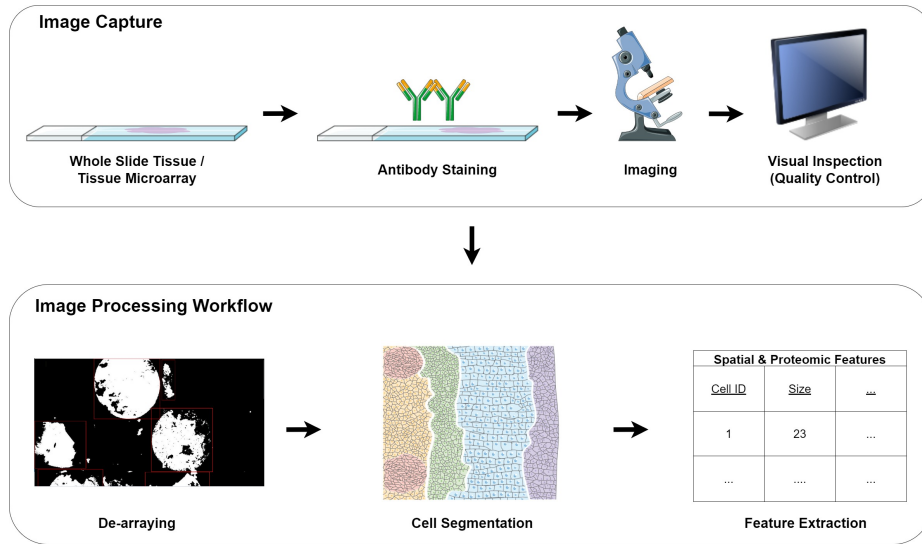


Fig. 7. An overview of the image capture and image processing stage for highly-plexed immunofluorescence imaging^a.

^aParts of the figure were drawn by using pictures from Servier Medical Art. Servier Medical Art by Servier is licensed under a Creative Commons Attribution 3.0 Unported License (<https://creativecommons.org/licenses/by/3.0/>).

porting the migration of current designs to CdB or to the Cinco-cloud platform that is currently being implemented.

Cinco de Bio uses semantic typing and data/process ontologies to tailor the concepts in its modelling language to the concepts and relations of the specific domain. This enables for example semantic compatibility checking of the data flow, which are in reality all TIFF files. Fig. 8 shows part of the data model taxonomy for the highly-plexed tissue image analysis A-DSL.

It also enables the reusability of domain-agnostic components, ensuring at the design stage that domain specific tools such as ML models (trained for a specific task) are being used correctly. The modelling language and the execution platform are designed to support both automated and interactive components: this is necessary because users must verify during analysis the outcome of workflow steps which are non-deterministic (such as for AI models) before proceeding to the following steps of the workflow.

In the context of our motivating use case, pre-processing highly-plexed tissue images for down-stream spatial and proteomic analysis there are a number of different segmentation algorithms developed to solve specific challenges. The first of which is known as de-arraying in the terminology of the field, which essentially entails cropping out each individual tissue core from a tissue micro array slide that may have over 100 cores. De-Arraying is done on a spatially

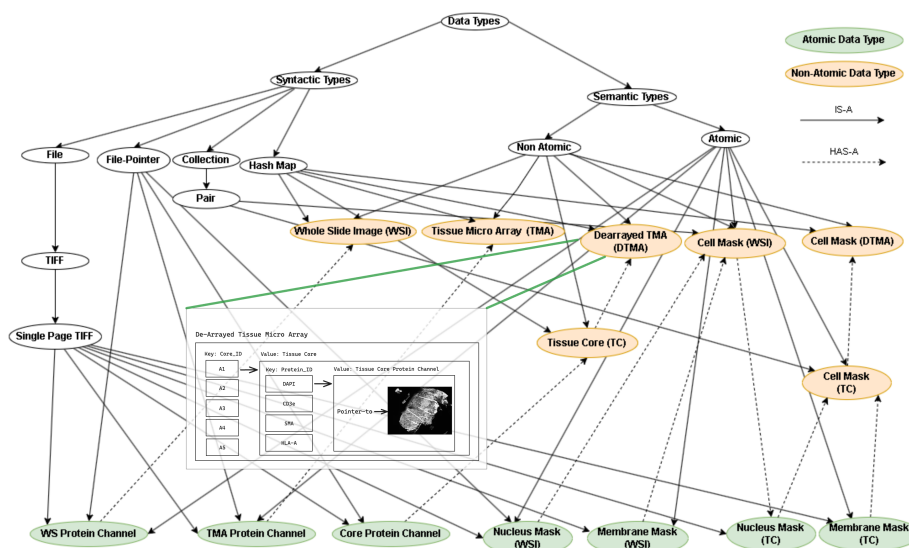


Fig. 8. An excerpt of the data model taxonomy for the highly-plexed tissue image analysis A-DSL. It classifies the application domain-specific data types in terms of their syntactic data types (as in computer science data structures and formats) as well as the semantic data types, expressing their meaning in the application domain. We distinguish also atomic and non-atomic data types. An atomic type is a single irreducible piece of data (in the context of the given application domain). For example, a Core Protein Channel equates to a single page tiff (greyscale image) file. A non-atomic type is a data structure which acts as a collection containing atomic or non-atomic types.

downsampled version of the image and the crop coordinates are then translated back to the full size images. The second form of segmentation is cell-segmentation which is done on the full resolution image (for each cropped core), to acquire the pixel coordinates for the nucleus and membrane of each cell. In both cases convolutional neural networks are typically used for the mask prediction task, however there are an array of post-processing steps required to acquire the final masks (especially in the context of cell segmentation).

As Cinco de Bio is the first case of the new platform architecture, we developed the modelling and execution platform, as well as the process/data ontologies and the concrete SIBS and workflow models. All the data-processing tools are implemented as containerised micro-services. This means that users can chain together heterogeneous data-processing services, that require different computational environments, into a single reproducible workflow.

The main contribution of this case study is the creation of a platform which prevents biologists from having to learn to code (here, Python and R) in order to undertake analyses on the data generated from their experiments. Secondly the

platform is designed for retargeting and extension to any domain, through the creation of new SIBs, workflows, and data and process ontologies which describe that domain.

4 Discussion

The four cases chosen for this contribution illustrate four different situations that are typical when advanced software engineering meets AI:

1. **MyMM** embraces MDD for the comfort of having an easily modifiable web application, but it does not do the step towards ML or AI. It would be possible to move to an ML implementation of the risk stratification, for example with decision trees or Algebraic Decision Diagrams (ADDs) in the Addlib [22], or with random forests, and achieve a possibly more performant computation. Here, however, the focus is on recognizability of the algorithm at the process and processing level, which limits what can be done and excludes approaches that would hide the structure of the computation, even if semantically equivalent.
2. **CNN-HAR** embraces ML and the data science lifecycle, but not the MDD: it contains classification and recognition, but is implemented essentially directly in Python. The workflow depicted would lend itself to the transformation to a Pyrus pipeline (if concerning just the data analysis) or a DIME application (if including more interaction with the IoT device and a web application). The limiting factor so far has been the choice of the *Thingy:53*, which is thought as a product for end users and not for integration: its own SDK cannot be accessed in a headless fashion, making the encapsulation into SIBs impossible or very onerous. Other devices like the *Thingy:52* are designed for designers and developers: it is a compact multi-sensor prototyping platform designed to help in building prototypes and demos, and it comes with a bluetooth Application Programming Interface (API) that is amenable to the transformation of commands into collections of SIBs.
3. **HIPPP** combines both the DIME-built web application with a heavy role of AI, ML and more, all encapsulated in reusable SIBs and processes. In this sense, HIPPP is a great example of what both disciplines can deliver when they join forces. HIPPP is a nice application, but being built on/with DIME it still requires local installation of the IME environment (Java, Eclipse etc), which is impractical for non-programmers who would otherwise not use these tools. It also does not support collaborative modelling, as the models reside in DIME and are not shared.
4. **Cinco de Bio** is the first prototype of the new generation of tools based on Cinco and soon on Cinco-cloud [7]. CdB address several limitations of the DIME application, namely CdB is a cloud native application whereas DIME is not. DIME has limitations with respect to the size/quantity of data it can process, therefore it cannot process data at the scale needed in the use-cases CdB is designed for. Finally incorporating non-Java based compute environments in DIME applications is a non-trivial task: the integration

could be handled with relative ease but the environments would have to be deployed separately and the address hardcoded into DIME.

On the other hand, CdB does not yet support hierarchical process modelling, the IME component is currently Cinco-based and has not been integrated to Cinco-Cloud and the CdB execution environment does not yet support GPU acceleration. These features are going to be addressed in future versions of CdB.

These observations are helping us develop the new R@ISE platform, which will run on Cinco-cloud and is heavily informed by the CdB experience.

In terms of collaboration, this has been a successful case study where we have started with a functioning implementation, partly in Python and partly in R, of the core software, provided by the cell biologists, who have learned Python and R, with the entire processing requiring a number of manual steps to transfer files, start pipelines, interact with the system, which was error prone and time consuming. Our task was to identify the core functionalities of the software, wrap them into individual SIBs, and reconnect them through workflows in an environment that then takes care of compilation and deployment (on heterogeneous runtime stacks, for Python and R), providing a unified web experience that automates all these steps and provides an easier experience as well as ease of reconfiguration and modification of the processes. This seems to be a task along the prior experiences with Bio-jETi [39], Fiat Flux [20], and ci:grasp [5, 6]

We had to intervene more than expected on the original code: disentangling it into separate, reusable SIBs brought to a complete refactoring of the code, and in that course we also optimised certain computations improving the readability of the code and the runtime performance.

We actually achieved more than this: the taxonomies are a new addition, the development of the interactive environment that supports human intervention for choices is new, and two new algorithms that address challenges in the application domain has also been devised. The first of which enables the correction of technical variances in samples. The other is a domain-specific data augmentation technique which enables the training of segmentation models on an extremely limited number of training samples. Both algorithms are now the subject of separate publications.

5 Conclusion

We have discussed four case studies from our recent research activity with real healthcare stakeholders, that illustrate different quadrants in the bidimensional space of advanced model driven development, in a low-code/no code fashion, and AI/ML.

In general, we have endeavoured to engage with the biology, health and medicine partners on the basis of their needs at that time, and their capabilities and interests. From previous experience, it would not have been helpful to try to impose a LC/NC approach on the HAR project at that time, because its focus

was on the efficient and effective edge computation capability, and not on generalizability. From that interaction, however, a collaboration grew, so that we are now developing further variants of the evaluation pipeline where our technology have found application.

On the contrary, the initial conversations with cell biologists several years ago concerned how to provide their new (AI and ML based) algorithms, that were implemented in Python, to third parties without handing out the Python code, which is an intellectual property embodiment they wished to protect. This led to conversations about interpreted vs. compiled languages: the distinction was new to them. They embraced the usefulness of abstraction (have a second artifact that is different from source code, for example compiled code, with different properties, for example executability and obfuscation) but they did not have the ability nor the funding for a reimplement in some compiled language. Years later we reconnected on the tissue analysis case study, where the ability to share the models and even give access to execution to third parties, at the same time hiding the source code was found very attractive and led to the design of Cinco de Bio.

While the healthcare partners were happy with what they got, and sometimes even surprised by what was achieved with the new technologies, it is clear to us that the future lies in the next step beyond Cinco de Bio. In the new R@ISE project, a large Strategic Partnership Project running for 5 years and co-funded by industrial and civil society partners and Science Foundation Ireland, we are currently starting to develop a new platform that will be online and cloud based, support the interaction and co-design directly on the models, support the semantic modelling through taxonomies as well as a strong capability to reason and evaluate the correctness of models at design time. We see here a role for fully automatic, "lightweight" formal methods [57] (static analysis, model checking, some synthesis) [42], but also for the use of AI, ML and Large Language Models (LLMs) in the platform itself, along the lines already recently initiated by partners in Dortmund [16] and [17]. Healthcare will be still a core application domain, in collaboration with the Bernal Institute, the Health Research Institute, the Limerick Digital Cancer Research Centre (LDCRC) and the University Hospital Limerick.

Acknowledgements

This work was conducted with the financial support of Science Foundation Ireland (SFI) under grants number 21/SPP/9979 (R@ISE), and 18/CRT/6223 574 (SFI Centre of Research Training in AI) as well as University of Limerick Health Research Institute ULCaN grant Pillar 4.

References

1. Gartner Forecasts Worldwide Low-Code Development Technologies Market to Grow 20% in 2023.

2. Minio inc: Minio — high performance, kubernetes native object storage. <https://min.io/>. Accessed: 2024-05-21.
3. Yousef Abuseta. A context aware framework for iot based healthcare monitoring systems. *International Journal of Advanced Studies in Computers, Science and Engineering*, 9(7):1–10, 2020.
4. Mugahed A Al-Antari, Mohammed A Al-Masni, and Tae-Seong Kim. Deep learning computer-aided diagnosis for breast lesion in digital mammogram. *Deep Learning in Medical Image Analysis: Challenges and Applications*, pages 59–72, 2020.
5. Samih Al-Areqi, Anna-Lena Lamprecht, and Tiziana Margaria. Constraints-driven automatic geospatial service composition: Workflows for the analysis of sea-level rise impacts. In *Computational Science and Its Applications–ICCSA 2016: 16th International Conference, Beijing, China, July 4–7, 2016, Proceedings, Part III 16*, pages 134–150. Springer, 2016.
6. Samih Al-areqi, Anna-Lena Lamprecht, Tiziana Margaria, Steffen Kriewald, Dominik Reusser, and Markus Wrobel. Agile workflows for climate impact risk assessment based on the ci: grasp platform and the jabc modeling framework. 2014.
7. Alexander Bainczyk, Daniel Busch, Marco Krumrey, Daniel Sami Mitwalli, Jonas Schürmann, Joel Tagoukeng Dongmo, and Bernhard Steffen. Cinco cloud: a holistic approach for web-based language-driven engineering. In *International Symposium on Leveraging Applications of Formal Methods*, pages 407–425. Springer, 2022.
8. Marco Bakera, Tiziana Margaria, Clemens D. Renner, and Bernhard Steffen. Tool-supported enhancement of diagnosis in model-driven verification. *Innovations in Systems and Software Engineering*, 5(3):211–228, September 2009.
9. Charos Boboyorova. An Introduction To No-Code/Low-Code Platforms: Seven Things You Should Know.
10. Luigi Borzi, Ivan Mazzetta, Alessandro Zampogna, Antonio Suppa, Gabriella Olmo, and Fernanda Irrera. Prediction of freezing of gait in parkinson’s disease using wearables and machine learning. *Sensors*, 21(2):614, 2021.
11. Luigi Borzi, Luis Sigcha, Daniel Rodriguez-Martín, and Gabriella Olmo. Real-time detection of freezing of gait in parkinson’s disease using multi-head convolutional neural networks and a single inertial sensor. *Artificial intelligence in medicine*, 135:102459, 2023.
12. Steve Boßelmann, Markus Frohme, Dawid Kopetzki, Michael Lybecait, Stefan Naujokat, Johannes Neubauer, Dominic Wirkner, Philip Zweihoff, and Bernhard Steffen. Dime: a programming-less modeling environment for web applications. In *Leveraging Applications of Formal Methods, Verification and Validation: Discussion, Dissemination, Applications: 7th International Symposium, ISoLA 2016, Imperial, Corfu, Greece, October 10–14, 2016, Proceedings, Part II 7*, pages 809–832. Springer, 2016.
13. Colm Brandon, Steve Boßelmann, Amandeep Singh, Stephen Ryan, Alexander Schieweck, Eanna Fennell, Bernhard Steffen, and Tiziana Margaria. Cinco de bio: A low-code platform for domain-specific workflows for biomedical research. *BioMedInformatics*, (to appear), 2024.
14. Colm Brandon, Adam J Doherty, Dervla Kelly, Desmond Leddin, and Tiziana Margaria. Hipp: Health information portal for patients and public. *Applied Sciences*, 13(16):9453, 2023.
15. Brendan Burns, Joe Beda, Kelsey Hightower, and Lachlan Evenson. *Kubernetes: up and running*. ” O’Reilly Media, Inc.”, 2022.
16. Daniel Busch, Alexander Bainczyk, and Bernhard Steffen. Towards llm-based system migration in language-driven engineering. In *International Conference on Engineering of Computer-Based Systems*, pages 191–200. Springer, 2023.

17. Daniel Busch, Gerrit Nolte, Alexander Bainczyk, and Bernhard Steffen. Chatgpt in the loop: a natural language extension for domain-specific modeling languages. In *International Conference on Bridging the Gap between AI and Reality*, pages 375–390. Springer, 2023.
18. Allan Fábio de Aguiar Barbosa. Model-driven development applied to mobile health and clinical scores. In *Evaluation of Novel Approaches to Software Engineering: 14th International Conference, ENASE 2019, Heraklion, Crete, Greece, May 4–5, 2019, Revised Selected Papers 14*, pages 188–203. Springer, 2020.
19. Florenc Demrozi, Graziano Pravadelli, Azra Bihorac, and Parisi Rashidi. Human activity recognition using inertial, physiological and environmental sensors: A comprehensive survey. *IEEE Access*, 8:210816–210836, 2020.
20. Birgitta E Ebert, Anna-Lena Lamprecht, Bernhard Steffen, and Lars M Blank. Flux-p: automating metabolic flux analysis. *Metabolites*, 2(4):872–890, 2012.
21. James P Gleeson, Thomas Brendan Murphy, Joseph D O’Brien, Nial Friel, Norma Bargary, and David JP O’Sullivan. Calibrating covid-19 susceptible-exposed-infected-removed models with time-varying effective contact rates. *Philosophical Transactions of the Royal Society A*, 380(2214):20210120, 2022.
22. Frederik Gossen, Alnis Murtovi, Philip Zweihoff, and Bernhard Steffen. Add-lib: Decision diagrams in practice. *arXiv preprint arXiv:1912.11308*, 2019.
23. Noah F Greenwald, Geneva Miller, Erick Moen, Alex Kong, Adam Kagel, Thomas Dougherty, Christine Camacho Fullaway, Brianna J McIntosh, Ke Xuan Leow, Morgan Sarah Schwartz, et al. Whole-cell segmentation of tissue images with human-level performance using large-scale data annotation and deep learning. *Nature biotechnology*, 40(4):555–565, 2022.
24. Ivan Guevara, Stephen Ryan, Amandeep Singh, Colm Brandon, and Tiziana Margaria. Edge iot prototyping using model-driven representations: A use case for smart agriculture. *Sensors*, 24(2):495, 2024.
25. Pandia Rajan Jeyaraj and Edward Rajan Samuel Nadar. Computer-assisted medical image classification for early diagnosis of oral cancer employing deep learning algorithm. *Journal of cancer research and clinical oncology*, 145:829–837, 2019.
26. Sven Jorges, Christian Kubczak, Felix Pageau, and Tiziana Margaria. Model Driven Design of Reliable Robot Control Programs Using the jABC. In *Fourth IEEE International Workshop on Engineering of Autonomic and Autonomous Systems (EASe’07)*, pages 137–148, March 2007.
27. John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Židek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.
28. Vaishnavi Kannan, Jason C Fish, and DuWayne L Willett. Agile model driven development of electronic health record-based specialty population registries. In *2016 IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)*, pages 465–468. IEEE, 2016.
29. Peter Kelb, Tiziana Margaria, Michael Mendler, and Claudia Gsottberger. Mosel: A flexible toolset for monadic second-order logic. In *Tools and Algorithms for the Construction and Analysis of Systems*, Lecture Notes in Computer Science, pages 183–202, Berlin, Heidelberg, 1997. Springer.
30. Jina Kim, Jieon Lee, Eunil Park, and Jinyoung Han. A deep learning model for detecting mental illness from user content on social media. *Scientific reports*, 10(1):11846, 2020.

31. Laura Kinkead, Ahmed Allam, and Michael Krauthammer. Autodiscern: rating the quality of online health information with hierarchical encoder attention-based neural networks. *BMC medical informatics and decision making*, 20(1):1–13, 2020.
32. Thomas Kluyver, Benjamin Ragan-Kelley, Fernando Pérez, Brian E Granger, Matthias Bussonnier, Jonathan Frederic, Kyle Kelley, Jessica B Hamrick, Jason Grout, Sylvain Corlay, et al. Jupyter notebooks—a publishing format for reproducible computational workflows. *Elpub*, 2016:87–90, 2016.
33. Saul A. Kripke. Semantical Analysis of Modal Logic I Normal Modal Propositional Calculi. *Mathematical Logic Quarterly*, 9(5-6):67–96, January 1963.
34. Ashnil Kumar, Jinman Kim, David Lyndon, Michael Fulham, and Dagan Feng. An ensemble of fine-tuned convolutional neural networks for medical image classification. *IEEE journal of biomedical and health informatics*, 21(1):31–40, 2016.
35. Farid Lahboube and Ounsa Roudies Nissrine Souissi. Building a his supervision metamodel. In *2016 11th System of Systems Engineering Conference (SoSE)*, pages 1–6. IEEE, 2016.
36. Anna-Lena Lamprecht, Tiziana Margaria, and Bernhard Steffen. Data-Flow Analysis as Model Checking Within the jABC. In *Compiler Construction*, Lecture Notes in Computer Science, pages 101–104, Berlin, Heidelberg, 2006. Springer.
37. Michael Y Lee, Jacob S Bedia, Salil S Bhate, Graham L Barlow, Darci Phillips, Wendy J Fantl, Garry P Nolan, and Christian M Schürch. Cellseg: a robust, pre-trained nucleus segmentation and pixel quantification software for highly multiplexed fluorescence images. *BMC bioinformatics*, 23(1):46, 2022.
38. Tiziana Margaria. Making Sense of Complex Applications: Constructive Design, Features, and Questions. In *Models, Mindsets, Meta: The What, the How, and the Why Not? Essays Dedicated to Bernhard Steffen on the Occasion of His 60th Birthday*, Lecture Notes in Computer Science, pages 129–148. Springer International Publishing, Cham, 2019.
39. Tiziana Margaria, Christian Kubczak, and Bernhard Steffen. Bio-jeti: a service integration, design, and provisioning platform for orchestrated bioinformatics processes. *BMC bioinformatics*, 9:1–17, 2008.
40. Tiziana Margaria, Daniel Meyer, Christian Kubczak, Malte Isberner, and Bernhard Steffen. Synthesizing Semantic Web Service Compositions with jMosel and Golog. In *The Semantic Web - ISWC 2009*, Lecture Notes in Computer Science, pages 392–407, Berlin, Heidelberg, 2009. Springer.
41. Tiziana Margaria and Bernhard Steffen. Backtracking-free design planning by automatic synthesis in metaframe. In *Fundamental Approaches to Software Engineering*, Lecture Notes in Computer Science, pages 188–204, Berlin, Heidelberg, 1998. Springer.
42. Tiziana Margaria and Bernhard Steffen. Backtracking-free design planning by automatic synthesis in metaframe. In *International Conference on Fundamental Approaches to Software Engineering*, pages 188–204. Springer, 1998.
43. Tiziana Margaria and Bernhard Steffen. Business Process Modelling in the jABC: The One-Thing-Approach. In Jorge Cardoso and Wil van der Aalst, editors, *Handbook of Research on Business Process Modeling*. IGI Global, 2009.
44. Tiziana Margaria and Bernhard Steffen. Service-Oriented: Conquering Complexity with XMDD. In Mike Hinchey and Lorcan Coyle, editors, *Conquering Complexity*, pages 217–236. Springer London, 2012.
45. Tiziana Margaria and Bernhard Steffen. eXtreme Model-Driven Development (XMDD) Technologies as a Hands-On Approach to Software Development Without Coding. In *Encyclopedia of Education and Information Technologies*, pages 732–750. Springer International Publishing, Cham, 2020.

46. Senthilkumar Mohan, Chandrasegar Thirumalai, and Gautam Srivastava. Effective heart disease prediction using hybrid machine learning techniques. *IEEE access*, 7:81542–81554, 2019.
47. MongoDB, Inc. Mongoddb: The developer data platform. <https://www.mongodb.com>, 2024. Accessed: 2024-05-21.
48. Ankit Murarka, Balaji Radhakrishnan, and Sushma Ravichandran. Classification of mental illnesses on social media using roberta. In *Proceedings of the 12th international workshop on health text mining and information analysis*, pages 59–68, 2021.
49. Shan Nan, Pieter Van Gorp, Hendrikus HM Korsten, Uzay Kaymak, Richard Vdovjak, Xudong Lu, and Huilong Duan. Dccss: A meta-model for dynamic clinical checklist support systems. In *2015 3rd International Conference on Model-Driven Engineering and Software Development (MODELSWARD)*, pages 272–279. IEEE, 2015.
50. Miguel Angel Olivero, FJ Domínguez-Mayo, Carlos Luis Parra-Calderón, MJ Escalona, and Alicia Martínez-García. Facilitating the design of hl7 domain models through a model-driven solution. *BMC medical informatics and decision making*, 20:1–18, 2020.
51. O Blake R Clifford and A Reche. Uhl guidance on management of mgus in primary care, 2019.
52. O Blake R Clifford and A Reche. Uhl guidance on management of mgus in primary care. Technical report, University Hospital Limerick, 2019.
53. Julie M Robillard, Jessica H Jun, Jen-Ai Lai, and Tanya L Feng. The quest for quality online health information: validation of a short quantitative tool. *BMC medical informatics and decision making*, 18:1–15, 2018.
54. Adam Sadilek, Henry Kautz, Lauren DiPrete, Brian Labus, Eric Portman, Jack Teitel, and Vincent Silenzio. Deploying nemesis: Preventing foodborne illness by data mining social media. *Ai Magazine*, 38(1):37–48, 2017.
55. Amandeep Singh, Tiziana Margaria, and Florenc Demrozi. Cnn-based human activity recognition on edge computing devices. In *2023 IEEE International Conference on Omni-layer Intelligent Systems (COINS)*, pages 1–4, 2023.
56. Bernhard Steffen, Andreas Claßen, Marion Klein, Jens Knoop, and Tiziana Margaria. The fixpoint-analysis machine. In *CONCUR '95: Concurrency Theory, Lecture Notes in Computer Science*, pages 72–87, Berlin, Heidelberg, 1995. Springer.
57. Bernhard Steffen and Tiziana Margaria. Tools get formal methods into practice. *ACM Computing Surveys (CSUR)*, 28(4es):126–es, 1996.
58. Dandan Tao, Ruofan Hu, Dongyu Zhang, Jasmine Laber, Anne Lapsley, Timothy Kwan, Liam Rathke, Elke Rundensteiner, and Hao Feng. A novel foodborne illness detection and web application tool based on social media. *Foods*, 12(14):2769, 2023.
59. Dandan Tao, Dongyu Zhang, Ruofan Hu, Elke Rundensteiner, and Hao Feng. Crowdsourcing and machine learning approaches for extracting entities indicating potential foodborne outbreaks from social media. *Scientific reports*, 11(1):21678, 2021.
60. Algimantas Venčkauskas, Vytautas Štuikys, Jevgenijus Toldinas, and Nerijus Jusas. A model-driven framework to develop personalized health monitoring. *Symmetry*, 8(7):65, 2016.
61. Sheng Wang, Zihao Zhao, Xi Ouyang, Qian Wang, and Dinggang Shen. Chatcad: Interactive computer-aided diagnosis on medical image using large language models. *arXiv preprint arXiv:2302.07257*, 2023.

62. Yu-Tseng Wang, Hen-Hsen Huang, Hsin-Hsi Chen, and H Chen. A neural network approach to early risk detection of depression and anorexia on social media text. In *CLEF (Working Notes)*, pages 1–8, 2018.
63. Pei-Ying Wen and Shien-Young Chang. Design and implementation of model-driven development for nursing information system. In *MEDINFO 2021: One World, One Health—Global Partnership for Digital Innovation: Proceedings of the 18th World Congress on Medical and Health Informatics*, volume 290, page 154. IOS Press, 2022.
64. Jason Williams. *RabbitMQ in action: distributed messaging for everyone*. Simon and Schuster, 2012.
65. Samir S Yadav and Shivajirao M Jadhav. Deep convolutional neural network based medical image classification for disease diagnosis. *Journal of Big data*, 6(1):1–18, 2019.
66. Amani Yahyaoui, Akhtar Jamil, Jawad Rasheed, and Mirsat Yesiltepe. A decision support system for diabetes prediction using machine learning and deep learning techniques. In *2019 1st International informatics and software engineering conference (UBMYK)*, pages 1–4. IEEE, 2019.
67. Qingda Zang, Kamel Mansouri, Antony J Williams, Richard S Judson, David G Allen, Warren M Casey, and Nicole C Kleinstreuer. In silico prediction of physicochemical properties of environmental chemicals using molecular fingerprints and machine learning. *Journal of chemical information and modeling*, 57(1):36–49, 2017.
68. Zizhao Zhang, Yuanpu Xie, Fuyong Xing, Mason McGough, and Lin Yang. Mdnnet: A semantically and visually interpretable medical image diagnosis network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6428–6436, 2017.
69. Philip Zweihoff and Bernhard Steffen. Pyrus: an online modeling environment for no-code data-analytics service composition. In *Leveraging Applications of Formal Methods, Verification and Validation: 10th International Symposium on Leveraging Applications of Formal Methods, ISoLA 2021, Rhodes, Greece, October 17–29, 2021, Proceedings 10*, pages 18–40. Springer, 2021.