TECHNICAL REPORTS

RESEARCH-DRIVEN APPROACH TO STANDARDIZATION

ICT, CONSTRUCTION AND AEROSPACE

Version 1.0 · October 2024
ISBN 978-99987-734-3-1

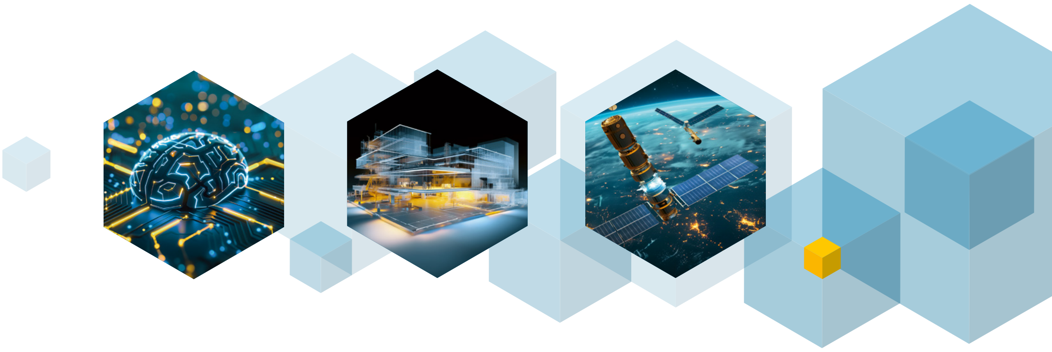

TECHNICAL REPORTS

RESEARCH-DRIVEN APPROACH TO STANDARDIZATION

ICT, CONSTRUCTION AND AEROSPACE

Luxembourg, October 2024

Institut Luxembourgeois de la Normalisation, de l'Accréditation, de la Sécurité et qualité des produits et services

Agence pour la Normalisation et l'Economie de la Connaissance

UNIVERSITÉ DU LUXEMBOURG

Acknowledgements

The working group (WG) involved in the preparation of these technical reports is:

Name of the contributor	Position	Institution/Organization
Prof. Dr. Pascal Bouvry	Dean of FSTM faculty	University of Luxembourg
Dr. Grégoire Danoy	Head of PCO group	University of Luxembourg
Dr. Mohammed Alswaitti	Postdoctoral Researcher	University of Luxembourg
Mr. Manuel Combarro Simón	Doctoral Researcher	University of Luxembourg
Ms. Hedieh Haddad	Doctoral Researcher	University of Luxembourg
Ms. Maria Hartmann	Doctoral Researcher	University of Luxembourg
Dr. Jean-Philippe Humbert	Deputy Director	ILNAS
Dr. Jean Lancrenon	Cybersecurity project officer	ILNAS
Ms. Victoria Mletzak	Responsible construction and Technical standardization	ILNAS
Dr. Lucas Cicero	Responsible aerospace and Technical standardization	ILNAS
Ms. Natalia Vinogradova	Head of Standardization Department	ANEC G.I.E.

This work is funded by the joint research program "Technical Standardization for Trustworthy ICT, Aerospace, and Construction (2021-2024)" between ILNAS and the Interdisciplinary Centre for Security, Reliability and Trust (SnT) of the University of Luxembourg.

Table of contents

Introduction

Technical standards are highly useful in laying the foundations of trustworthiness in products, processes, and services across all kinds of domains. The reason for this lies essentially in their objective and their elaboration process. Indeed, standards aim to yield broad, high-quality, flexible solutions to general problems encountered in industry, commerce, and society while placing as many stakeholders as possible around the table to reach consensus on the adopted solutions. These principles, while non-exhaustive, are a part of the backbone of the technical standardization process, enabling the publication of guidelines, recommendations, and requirements that benefit safety, quality, innovation, and progress.

Cutting edge scientific research results also feed into this overall process. Indeed, technical advancement occurs via research by bringing theoretical fundamental concepts to life through the construction of prototypes and ultimately the production of practical applications. Hence, creating links between research and technical standardization has the potential to feed both communities for mutual benefit.

The ILNAS-UL research programs - defined and implemented jointly by ILNAS and the University of Luxembourg - have been active in this endeavor since 2018, in synchronization with Luxembourg's successive National Standardization Strategies. Accordingly, the current ILNAS-UL research program, *Technical Standardization for Trustworthy ICT, Aerospace and Construction 2021-2024*, is articulated around the three major economic sectors identified in the National Standardization Strategy 2020-2030, those being the ICT, Aerospace, and Construction sectors. More specifically, the program is being driven forward by a dedicated team of three PhD students who, in addition to their scientific work, are exploring how standards are relevant to their research and vice versa. One of the outcomes of this program was the 2023 white paper *Trustworthiness in ICT, Aerospace, and Construction applications - Scientific Research and Technical Standardization - October 2023* in which general directions were identified to illustrate how technical standardization and research could contribute together to enhance trustworthiness.

With the current publication, the students each take this work further and elaborate on the identified challenges based on the outcomes of their research.

In the first chapter, dedicated to ICT, a gap between standardization and research with respect to the creation of geographic imagery mosaics is identified and a way towards addressing this gap is investigated covering different aspects and challenges related to mosaic production. The second chapter, more focused on the Aerospace sector, explores the current state of the art for the communication of machine learning models in a space context and proposes a path towards their smooth integration in general, and innovative distributed learning solutions in particular, into standardized space communication protocols. Last but not least, in the third chapter, the emphasis is on the Construction sector, where it is shown how using technical standards as base input to modern ICT tools can lead to near-direct standards compliance in AI-generated, optimized solutions to complex construction design problems.

This joint effort illustrates the variety of synergies that can exist between technical standardization and research, with a view towards encouraging the national market to both get involved in the standards drafting process and contribute to solidifying trustworthiness in the systems on which we all depend.

Finally, it is worth noting that as of August 2024, the National Standardization Strategy has been updated for the 2024-2030 period, namely through 1) the addition of the domain of Sustainability and 2) the placement of a new emphasis on the domain of Conformity Assessment. Indeed, these are key transversal domains gaining ever more importance: the former to evolve towards a more environmentally respectful society and the latter to establish assurance of met requirements. Accordingly, a potential new iteration of an "ILNAS-UL research program (2025-2028)" will have to address the opportunity to tackle these new dimensions of the National Standardization Strategy.

1.

Extending geographic information imagery standards for mosaic generation

1. Extending geographic information imagery standards for mosaic generation

Abstract

Advances in space technologies have significantly improved the availability of satellite imagery, which is vital for applications such as disaster management, environmental monitoring, and crop assessment. These applications often require the analysis of extensive regions, necessitating the use of multiple images combined into a mosaic. Ideally, a mosaic should appear seamless, but achieving this requires complex processes like radiometric normalization and optimal seamline detection, which present significant challenges.

In some cases, it is advantageous to use images from multiple sources to create a mosaic, as a single satellite or constellation may not fully cover the desired area of interest. This solution increases flexibility and can enhance characteristics such as cost, cloud coverage, and image resolution. However, it introduces additional complexity, particularly due to varying metadata schemas and the difficulties of radiometric normalization across different sources. Further complications arise when users need to optimize multiple mosaic parameters simultaneously, such as generating the optimal set of solutions, known as the Pareto front.

This report examines these challenges, identifies the relevant geographic imagery standards and highlights the gaps related to mosaic generation. We present specific areas where these standards can be improved to ensure the creation of high-quality and reliable mosaics.

1.1. Introduction

Understanding and monitoring the Earth's surface is crucial for addressing a wide range of global challenges, from food security to environmental sustainability. Geographic imagery serves as a vital tool by providing a visual representation of the Earth's surface, offering a view from above and supporting environmental studies, analysis and forecasts. In recent years, with the progress in optical sensors and space technologies, satellite imagery availability has significantly increased, creating new opportunities for a large variety of applications. Some of these can rely on a single image, while others need to combine multiple images to cover extensive areas and meet the application's needs, such as disaster management or tracking of deforestation over time. However, obtaining the merged image, known as a mosaic, is not a straightforward process and a significant number of challenges remain to be overcome.

The most common way of creating mosaics is to combine images from a single satellite or satellite constellation. Nevertheless, in some cases, the quality of some of the images may hinder the exploitation of the final mosaic. Moreover, when considering the tracking of a certain area over time, the same satellite or satellite constellation may not have all the necessary images available to cover the zone of interest, bringing the need to combine images from different satellite constellations. These scenarios raise new challenges, such as the presence of a large number of potential mosaics covering the same area, which makes selecting the most appropriate one a complex task.

This report focuses on these different challenges and creates a link with the existing normative documents, where possible. Indeed, while numerous standards have been established to ensure the quality, interoperability, and reliability of geographic imagery, there remain notable gaps in standardization, particularly concerning mosaic generation. Thus, one of the objectives of the current document is to spotlight areas where standardization efforts could be improved to better support the needs of satellite image mosaic users.

1.2. Geographic imagery

Geographic imagery is defined in the ISO 19101-2:2018 - *Geographic information - Reference model - Part 2: Imagery*[1] standard as a "representation of phenomena as images produced by electronic and/or optical techniques associated with a location relative to the Earth". The two primary sources of geographic imagery are aerial and satellite imagery. Aerial imagery is captured from aircraft, such as balloons, planes, and drones, and covers smaller areas compared to satellite imagery, which offers more extensive coverage and has gained popularity due to its increased availability thanks to the *New Space* approach (more frequent satellite launches with lower related costs).

To set the scene and detail these challenges in the next sections, the steps inherent to the satellite image exploitation are considered, starting from their acquisition to the delivery of images to the user. This process is well described in the ISO 19101-2:2018 standard where it is modelled as a system of five interconnected nodes, as illustrated in Figure 1. These nodes are:

1) **Imagery Collection.** This node contains all the aspects related to the imaging sensor located in the satellite and oversees the image acquisition.
2) **Sensor Processing.** It is the node that communicates with the sensor and handles initial processing steps such as image compression and coordinate conversion.
3) **Imagery Archive.** Responsible for data management, including storage, access, preservation, and retrieval.
4) **Value Added Processing.** Involves additional processing and data manipulation to enhance its utility for specific applications.
5) **Decision Support:** Facilitates human interaction by providing imagery and geographic information to support decision-making processes. It gives access to both the *Imagery Archive* and *Value Added Processing* nodes. This node can utilize analytical methods, such as optimization algorithms, to help decision-makers select the most appropriate options for their applications.

Image mosaicking is one of the processing functions included within the *Value Added Processing and Decision Support* nodes.

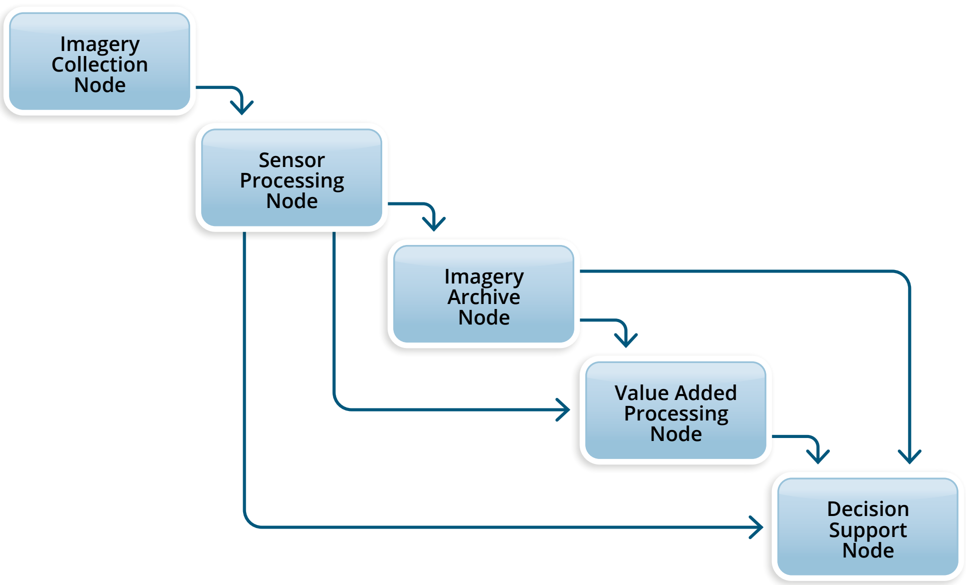

Figure 1. Geographic imagery system deployment diagram.

1 https://www.iso.org/standard/69325.html

1.3. Image mosaic

An image mosaic is defined in the OGC Image Exploitation Services[2] as "an assembly of two or more overlapping or adjacent orthorectified (or rectified) images to form a continuous image of a larger ground area". This process is executed when analyzing a large geographical area that cannot be covered by a single image, by merging several images to ensure full coverage of the area of interest. Figure 2 provides a visual demonstration of the need for the mosaic's creation.

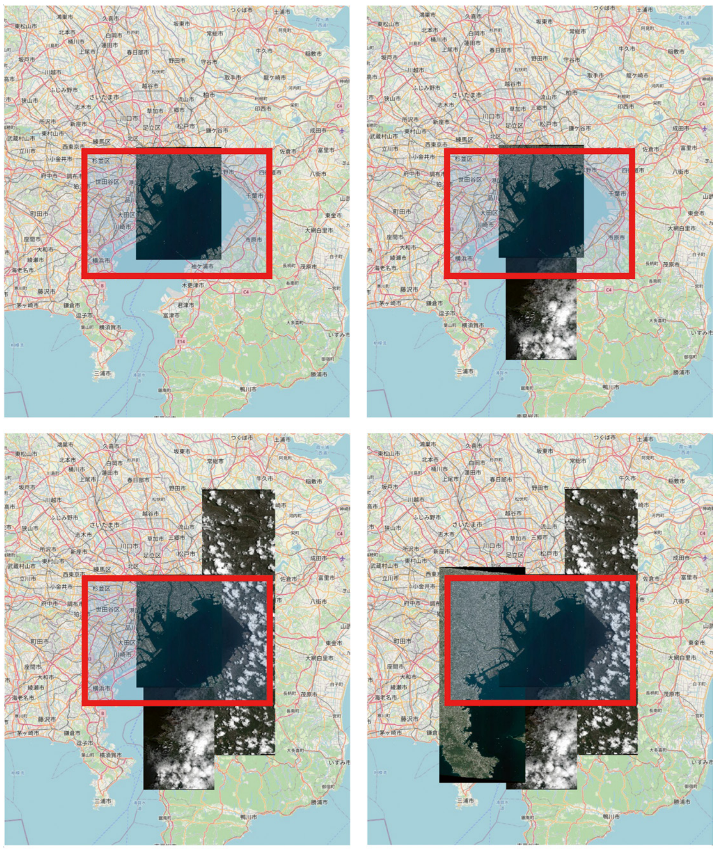

Figure 2. The area of interest (represented by the red rectangle) cannot be covered by only one image. A mosaic is needed to study the whole area of interest.

Following the definition, one of the important challenges in mosaic creation is the elimination of seamlines, so that the transition between individual images is imperceptible, and the final image looks uniform and continuous, with no visible borders between adjacent images. One simple scenario consists of creating the mosaic with axis-parallel rectangle images. This case is described in the ISO/AWI 19123-4 - *Geographic information - Schema for coverage geometry and functions - Part 4: Tiling Schema standard*[3] currently drafted by the technical committee ISO/TC 211 - *Geographic information/Geomatics*[4], where the rectangle images are referred to as "tiles". This future standard limits the mosaic creation to the use of axis-parallel rectangle images, with a low consideration about overlapping. However, in practical applications, images used for mosaicking may not always be axis-parallel rectangles and there will often be overlapping between them. Furthermore, overlapping is necessary in most cases to create seamless mosaics since it allows for a correction of radiometric differences and an optimal placement of seamlines.

2 OpenGeospatial® Consortium Abstract Specification, Topic 15 — Image Exploitation Services, 2000, OGC document 00-115

3 https://www.iso.org/standard/88297.html

4 https://www.iso.org/committee/54904.htmlse

Referring back to the example depicted by Figure 2, the mosaic was created by merging a selection of images together without any additional processing. As a result, the borders of the images are clearly visible. To create seamless image mosaics with a smooth and coherent visual rendering, several key steps must be undertaken after selecting the images to cover the whole area of interest:

- **Radiometric normalization (tonal adjustment):** This process adjusts the brightness and color levels of images to ensure uniformity across the mosaic. It simulates consistent illumination conditions, even if the images were captured under varying lighting conditions. This step is crucial when images are captured at different times because variations in radiometric values can lead to noticeable discrepancies.
- **Seamline detection and image blending:** Identifying the optimal boundaries or "seams" for blending overlapping images is critical. The goal is to place these seams in areas where transitions between images are the least noticeable. The final step aims to minimize any remaining differences along the seams. Techniques such as feathering or gradient blending are employed to smooth the transitions and unify the mosaic, this last step is beyond the scope of this report and will not be discussed here.

Radiometric normalization

Radiometric normalization of the images composing a mosaic starts with the identification of one image as the "source" and the others as the "targets". The adjustment of the radiometric properties of the target images is then made using a mapping function to match colorimetric characteristics of the source image. This function is derived from a comparative analysis of the overlapping region where both images cover the same geographical area, and then applied across the entire target image to ensure consistency, as illustrated in Figure 3.

According to [1] the techniques for radiometric normalization can be grouped into:

- **Global models:** These use a single mapping function to correct the target images, suitable for homogeneous terrains or low-resolution images.
- **Local models:** These use multiple mapping functions to address varying radiometric conditions within distinct parts of the image, which can be more obvious in high-resolution images or complex terrains. Local models must be applied carefully, as they can introduce new radiometric differences.
- **Combined models:** These integrate the global model and the local models, taking the best of both approaches. This method is particularly useful for high-spatial-resolution images and is considered one of the most promising techniques.

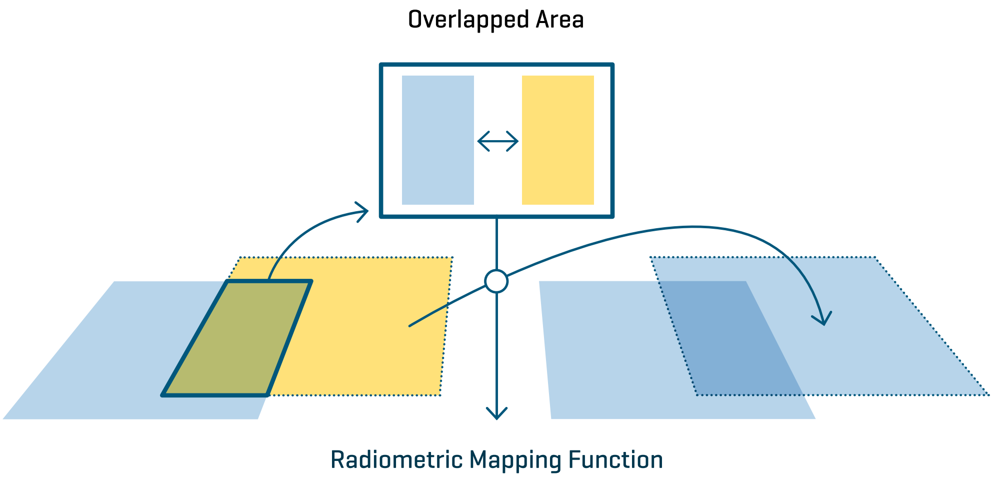

Figure 3. Radiometric normalization. The radiometric values of the target image (yellow rectangle) are normalized with respect to the source image (blue rectangle) using a mapping function obtained from the overlapping area. Based on [1].

Seamline detection

Once radiometric normalization is applied to the images, it is necessary to determine which part of the images will appear in the overlapping region of adjacent images. This involves detecting a seamline that divides the overlapping area, as seen in Figure 5. The selected seamline significantly influences the final mosaic's visual quality, as shown in Figure 4, where a non-optimal seamline was selected. The optimal seamline is located where geometric and pixel intensity differences between overlapping images are minimal [2].

Figure 4. Consequences of selecting a non-optimal seamline during the merging process. The selected seamline traverses the crop parcels, resulting in noticeable differences due to images captured at different stages of the crop cycle. In this case, a better seamline would have followed the road, to avoid crossing any crop parcel. Image taken from Google Maps[5].

According to [1] the techniques for optimal seamline detection can be categorized into:

- **Image internal information-based methods:** These rely solely on information from the images themselves. This category includes:
  - **Frame-to-frame methods:** Seamlines are detected sequentially for each pair of overlapping images.
  - **Multiframe joint methods:** All the seamlines are detected simultaneously.
- **External data-based methods:** These methods use additional information, such as the digital surface model (DSM), to improve seamline detection.

Considerations for standardization

Despite the complexity of these two steps, there are some guidelines that can be applied to facilitate these processes and obtain high-quality mosaics. In the following paragraphs, we will discuss these ones, while providing the relevant standards that support these tasks and highlighting the points not yet fully covered by the current standardization publications.

Overlap

Although the future ISO/AWI 19123-4 standard addresses partially images overlap, this scenario seems not to be considered as a mandatory requirement. However, some degree of overlap is necessary for effective radiometric normalization and seamline detection. Without overlap, generating a mapping function for radiometric normalization is impossible, and seamlines are restricted to the image edges, which are easily noticeable, especially in high resolutions.

5 https://rb.gy/i4vf7c

Additionally, if the overlapping area is too small, the mapping function may be biased, leading to unnatural tone balancing, and the quality of the mosaic may not be as high as it would be with a larger overlapping area. On the other hand, excessive overlap can be inefficient, as it uses only a small fraction of the images in the mosaic, potentially increasing costs. Therefore, defining a good range for overlap is crucial to balance quality and cost in mosaic generation.

Radiometric normalization

Beyond overlap, other considerations should be supported by standardization, such as the choice of radiometric normalization techniques: the selection of a radiometric normalization model should be guided by standardization, by considering external parameters, such as the resolution of the images and the complexity of the terrain. Global models may suffice for low-resolution images or homogeneous areas, while local or combined models are preferable for high-resolution images and complex terrains. It is worth mentioning that research done in the field of radiometric normalization, especially regarding the radiometric mapping function [3], has already employed tools defined by standards, such as the model color space definition presented in ISO/CIE 11664-4:2019 *Colorimetry - Part 4: CIE 1976 L*a*b* colour space*[6]. This should be included in a comprehensive standardized guideline as suggested.

To ease radiometric mapping function execution, consideration should be given to image metadata and the related colorimetry information. Indeed, standardization already offers methodologies and processes to quantify and harmonize the color balancing of optical sensors and which can directly support uniform mosaic creation. For example, the usage of ISO 19264-1:2021 - *Photography - Archiving systems - Imaging systems quality analysis - Part 1: Reflective originals*[7] along with ISO 14524:2009 - *Photography - Electronic still-picture cameras - Methods for measuring optoelectronic conversion functions (OECFs)*[8] by satellite manufacturers, can support the integration of tonal curve adjustment to image metadata. The ISO 19115 series defines the metadata schema of geographic imagery but currently does not include any information related to image colorimetry. A suggestion is made to add such information following a standardized format based, for example, on ISO 18620:2016 - *Graphic technology - Prepress data exchange - Tone adjustment curves exchange*[9].

Seamline detection

For seamline detection, the choice between frame-to-frame and multi-frame joint methods depends on the number of images and the mosaic's complexity. Frame-to-frame methods should be used only when the number of images is small [1]. Additionally, it is important to clearly define the optimal seamline and under what conditions it is necessary to detect it. Below we present relevant considerations related to the optimal seamline detection that can be included in future standards:

- For low and medium-resolution images, finding the optimal seamline is not generally necessary. The initial seamline can be a straight line that divides the overlapping region in half, and if necessary, it can be optimized later [1]. See Figure 5 (a) and (b).
- For high-resolution images (more detailed and smaller than low and medium-resolution), finding the optimal seamline is critical. This is because the difference in illumination is more notable, and as the images are small and highly detailed, moving objects can be seen, and the seamline should not cross them [4].
- The seamlines should be curved because the human eye can detect the seam better if it is a straight line [5]. See Figure 5 (c).

6 https://www.iso.org/standard/74166.html
7 https://www.iso.org/standard/79172.html
8 https://www.iso.org/standard/43527.html
9 https://www.iso.org/standard/63034.html

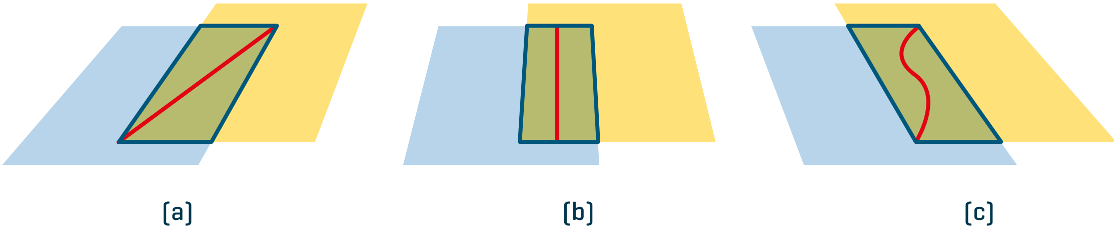

Figure 5. Example of seamlines. Cases (a) and (b) are used in low and medium-resolution images. (c) A curved seamline is more difficult for the human eye to detect. Based on [1].

1.4. Multi-source image mosaic

In both scientific literature and existing standards, image mosaics are typically considered to be created using images from the same source. This approach is commonly used when it is necessary to study an area in the future, such as monitoring the recovery of a region after a natural disaster for several days or months. In such cases, the user can request a satellite provider to capture images that cover the area of interest, resulting in a mosaic composed of on demand images taken by the same satellite or satellite constellation.

However, there are scenarios where a user might be interested in studying a large region's past by using archive images, such as environmental monitoring, agricultural trends or mapping. In these situations, the user does not have control over the available images, and there exists a risk that a single provider may not have enough archive images to cover the entire area of interest. This can be illustrated with a hypothetical situation where a user is interested in studying a vast region surrounding Luxembourg on the specific date of 15/07/2024. By consulting the EOSDA LandViewer[10] marketplace, which can only create mosaics using images from a single satellite constellation, the entire area of interest cannot be fully covered, highlighting the limitation of relying on a single satellite source (see Figure 6).

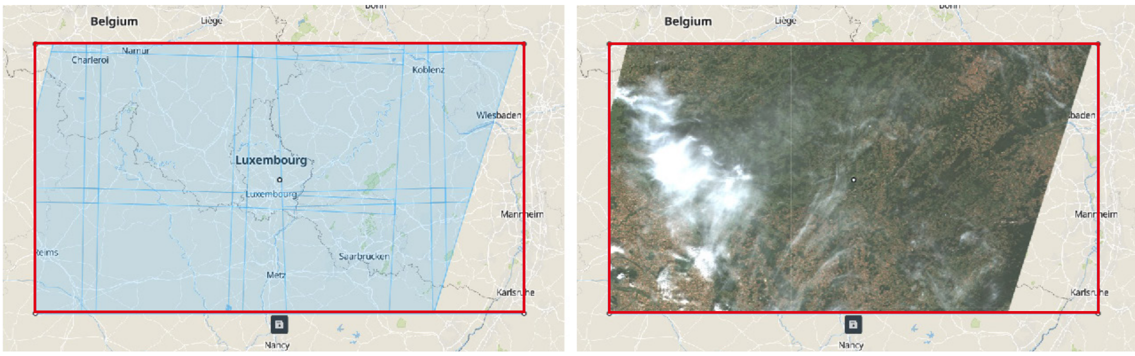

Figure 6. An area of interest, represented by a red edge rectangle, cannot be covered by a mosaic built using only images from Sentinel-2 for 15/07/2024 (blue polygons). The top-left and the right corners of the area of interest are not covered by the mosaic. The image on the left shows the overlap of the 15 images used to build the mosaic, while the image on the right displays the final mosaic. Image taken from EOSDA LandViewer.

10 https://eos.com/products/landviewer/

Mosaicking satellite images from different providers offers several advantages. First, as described previously, this approach may provide full spatial coverage, which is not always possible when using images from a single provider. Indeed, the restrictions on coverage or temporal sampling can differ for each provider, enabling the possibility to combine images from different sources to satisfy spatial requirements. With more available images, this process also improves the temporal dimension, which is very important in situations where changes are frequent and require constant updates. Additionally, using images from different sources provides more flexibility and can enhance the final mosaic regarding parameters such as cloud coverage or cost. For instance, a smart combination of images from different sources can reduce the cloud coverage percentage in the final image. In Figure 6, for example, an image captured by another satellite constellation on the same day – but at a different time – could cover the cloudy area on the left side of the picture, helping to create a cloudless mosaic. A similar approach can also optimize the cost of the mosaic, which is equal to the sum of the images' costs. Since each satellite image provider has a different pricing policy, it is possible to find a combination of images that reduces the total cost while meeting user requirements.

Moreover, when parameters need to be within certain ranges, such as spatial resolution (see Table 1), incidence angle or date interval of the images, combining images from different sources can be beneficial, as users often prefer that most of the images have the best possible values within defined ranges.

Considerations for standardization

Overall, combining images from multiple sources significantly enhances the flexibility and quality of image mosaics, addressing various user requirements more effectively. Nevertheless, the process of constructing multisource mosaics also reveals a critical gap in existing standards. Although there are some standards such as ISO 19101-2:2018 - *Geographic information - Reference model - Part 2: Imagery*[11] and ISO 19115-1:2014 - *Geographic information - Metadata - Part 1: Fundamentals*[12] that discuss some aspects of image processing and metadata schemas that can be used to integrate images from different sources, no clear guidelines are provided to fuse multisource data. Also, in the current standardization landscape, there are no well-defined procedures for the aligning, calibrating, and radiometric normalization of images acquired from different satellite constellations which poses many challenges towards a seamless multisource-mosaic creation process. Hence, there is a great need to come up with better and more effective standards that can efficiently deal with the complexities of multi-source satellite image mosaics ensuring their consistency, quality, and interoperability. Such standards would facilitate reliable integration of multi-sourced data and drive the advancement of satellite imagery and its applications. More globally, geographic imagery future standards should consider the possibility of creating mosaics from diverse sources and provide recommendations to address the associated challenges.

Scales of applications in urban areas		Image data used for these applications	
Applications	**Scale**	**Image Source**	**Resolution**
Facility management	1:200 to 1:500	Aerial photo/LIDAR	< 20 cm
Basic mapping	1:1000 to 1:2000	Aerial photo/LIDAR	20 cm to 50 cm
Urban planning	1:5000 to 1:10000	High-resolution satellite / Aerial photo	50 cm to 10 m
Overview	1:10000 to 1:1000000	Satellite	10 m to 30 m

Table 1. Examples of applications and image spatial resolution. Table taken from ISO ISO 19101-2:2018.

11 https://www.iso.org/standard/69325.html

12 https://www.iso.org/standard/53798.html

1.5. Multi-objective mosaic generation

In the previous section, we observed that mosaics, beyond simply covering large areas, can be constructed in ways that optimize certain parameters, such as reducing cloud coverage. However, in real situations, the user is interested in optimizing not one, but multiple parameters simultaneously. For example, high-resolution satellite images are very expensive, so cost is always a parameter to consider alongside the quality of the mosaic.

When users aim to optimize several parameters or objectives simultaneously, the process is known as multi-objective optimization. When the objectives are conflicting, it is unlikely that a single solution will achieve the best values for all objectives, as optimizing one objective often leads to the degradation of others. For example, in the case of satellite imagery, for similar images, the cost can be affected by the cloud coverage percentage; the lower this value the higher the cost of the image, so a mosaic with less cloud coverage will cost more. In such scenarios, where all objectives are equally important, a dominance relation between solutions can be established. A solution X is said to dominate another solution Y if X is better than Y in all objectives. If neither solution dominates the other across all objectives, they are considered non-dominated solutions.

In multi-objective problems, there exists one particular subset of non-dominated solutions, known as the Pareto front, which satisfies the following: any solution outside this subset is dominated by at least one solution from the subset. In Figure 7, the Pareto front for a bi-objective minimization problem is shown in yellow. We can observe that any solution outside the front is dominated by at least one solution from the front, while all solutions on the front are non-dominated. The goal of multi-objective optimization is to identify and present this subset of solutions, allowing decision-makers to choose a solution from the Pareto front based on their specific priorities and preferences.

To illustrate, consider three mosaics covering the same area: A, B and C, with costs of 100, 500 and 700 and cloud coverages of 50%, 20%, and 70%, respectively. If the primary objective is to minimize cost, mosaic A would be the preferred choice. However, if minimizing cloud coverage is equally important, mosaic C can be discarded as it is dominated by both A and B in terms of cost and cloud coverage. Mosaic B, despite being more expensive than A, offers significantly less cloud coverage, making both A and B nondominated solutions. Consequently, the system would present both A and B as potential options, leaving the final decision to the user based on their specific needs and preferences.

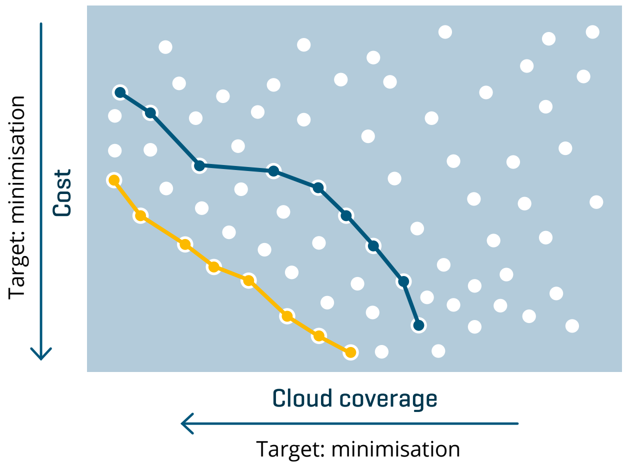

Figure 7. All the solutions of a minimization multi-objective mosaic generation problem with two objectives, cloud coverage and cost. Each point represents a solution, a possible mosaic. The Pareto front is shown in yellow, and an approximation is shown in blue. Any solution outside the Pareto front is dominated by at least one solution in the Pareto front.

In practice, multi-objective mosaic generation is far more complex than the earlier example: in fact, it has been proven to be a very complex computational problem [6]. For this problem, finding the Pareto front might take a significant amount of time unless the number of available images to build the mosaic is small or the user is interested in optimizing few objectives. If a fast response is needed, it is better to find an approximation front instead of the Pareto front, as illustrated by the blue front in Figure 7.

If users prioritize speed in decision-making, the exact methods should be used only when the number of objectives is small (e.g., two or three) and the dataset is small (e.g., 30 images). For larger problems or when multiple objectives are involved, non-exact methods may be more appropriate, providing approximations instead of the exact front.

An alternative way to improve the response time is to identify a representative subset, rather than computing the entire front. This can be also beneficial in cases where the front is too large. Suppose the front contains 1000 mosaics; this could be overwhelming for the user, making the decision to select the appropriate mosaic tedious. In such cases, it is better to return a limited number of mosaics, e.g. 50. Moreover, this approach can also be interactive, allowing users to initially receive a smaller set of solutions and then request additional solutions located in specific areas of the solution space.

Considerations for standardization

Users often desire to optimize more than one parameter when creating a mosaic, as not only one parameter is important for their application. For example, while using high-resolution images might be important for most of application, minimizing the cost of the mosaic is also desirable to maximize the profit.

As multi-objective mosaic generation is a very complex process, it is crucial to establish guidelines or standards for the characteristics of images used to create smooth mosaics, which may include adding new metadata fields for these characteristics. For instance, it should be clarified whether radiometric normalization can be applied to any pair of images, regardless of their intensity differences. This is essential because, during the problem-solving phase, no preliminary steps are taken to enhance the mosaic, and it is assumed that any combination of images will produce a seamless mosaic. If the system fails to guarantee that all proposed solutions will produce a seamless mosaic, the time spent to generate these faulty solutions is wasted, and more importantly, user confidence in the system will be undermined.

Furthermore, existing standards may need revision to include the possibility of multiple mosaics being generated from the available set of images as a result of multi-objective optimization. For instance, in ISO 19101-2:2018, this situation is not contemplated, as it is assumed that there is only one available mosaic, that can be generated in the *Value Added Processing* or *Decision Support* nodes from Figure 1. In the *Value Added Processing* node, mosaic creation may be limited to tasking mosaics, where the provider will capture all the necessary images to build a mosaic that covers the area of interest. However, when mosaics are created from archive images, especially in the context of multi-source mosaics, this process should occur in the *Decision Support* node, as multiple mosaics with different parameter values (such as cost, cloud coverage percentage, etc.) may need to be considered.

In such cases, the *Decision Support* node should implement methods for obtaining the Pareto front or an approximation, depending on the user's time constraints, the number of parameters to optimize, and the number of available images for mosaic creation. Additionally, the system should consider user preferences, such as specifying a desire for a limited number of solutions (e.g., K points or fewer) to improve the speed of solution generation.

1.6. Conclusion

Despite the existence of numerous standards for geographic imagery, significant gaps remain in addressing the complexities of mosaic generation. Current standards, such as those outlined in this report, do not fully consider the challenges involved in creating seamless mosaics, using images from multiple sources, or optimizing for multiple objectives simultaneously for mosaic generation.

There is a particular gap in standards regarding the creation of mosaics from diverse image sources. As users increasingly require historical data or seek to combine images from different satellite constellations, it becomes essential to establish common metadata schemes and provide guides to handle radiometric differences and metadata inconsistencies across sources.

Furthermore, the challenge of multi-objective optimization in mosaic generation is not adequately addressed. Users often desire to optimize, at the same time, conflicting parameters such as cost, resolution, and cloud coverage. Current standards do not provide sufficient guidance on how to handle scenarios where multiple satisfactory solutions exist, such as utilizing Pareto fronts to present to the users a range of optimal choices based on their preferences.

To enhance the utility and reliability of geographic imagery mosaics, future standards should incorporate detailed guidelines for seamless mosaic creation, address the complexities of multi-source mosaics, and provide frameworks for multi-objective optimization. By filling these gaps, standards can better support the diverse needs of users, ensuring high-quality, reliable, and efficient use of geographic imagery data.

References

[1] Li, X., Feng, R., Guan, X., Shen, H., Zhang, L.: Remote Sensing Image Mosaicking: Achievements and Challenges. IEEE Geosci. Remote Sens. Mag. 7, 8–22 (2019). https://doi.org/10.1109/MGRS.2019.2921780.

[2] Duplaquet, M.-L.: Building large image mosaics with invisible seam lines. Visual Information Processing VII. pp. 369–377. SPIE (1998). https://doi.org/10.1117/12.316427.

[3] Cresson, R., Saint-Geours, N.: Natural Color Satellite Image Mosaicking Using Quadratic Programming in Decorrelated Color Space. IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens. 8, 4151–4162 (2015). https://doi.org/10.1109/JSTARS.2015.2449233.

[4] Zagrouba, E., Barhoumi, W., Amri, S.: An efficient image-mosaicing method based on multifeature matching. Mach. Vis. Appl. 20, 139–162 (2009). https://doi.org/10.1007/s00138-007-0114-y.

[5] Fernandez, E., Garfinkel, R., Arbiol, R.: Mosaicking of Aerial Photographic Maps Via Seams Defined by Bottleneck Shortest Paths. Operations Research 46, 293–304 (1998). http://dx.doi.org/10.1287/opre.46.3.293.

[6] Combarro Simón, M., Talbot, P., Danoy, G., Musial, J., Alswaitti, M., Bouvry, P.: Constraint Model for the Satellite Image Mosaic Selection Problem. Leibniz Int. Proc. Informatics, LIPIcs. 280, 1–23 (2023). https://doi.org/10.4230/LIPIcs.CP.2023.44.

2. Extending communication standards to facilitate machine learning on board spacecraft

2. Extending communication standards to facilitate machine learning on board spacecraft

Abstract

In a time of rapid technological change, standardization is a crucial tool to facilitate the safe and effective deployment of cutting-edge technologies to the market. This is particularly true for the aerospace domain, where an emphasis on rigorous standardization is born of hard-earned experience. Nowadays, the aerospace domain is facing new technological opportunities from several fields at once: on one side, fundamental mission design concepts are undergoing a paradigm shift, from the historical monolithic satellite design to the new concept of distributed formations made up of many smaller satellites. On the other side, the recent advancements in artificial intelligence research have opened the door for a new level of task automation on-board spacecraft.

We explore how standardization could lay the groundwork for a type of space mission that arises at the intersection of these two new developments: the on-board deployment of machine learning models across multiple small spacecraft. In particular, we examine the challenges of communicating such models from ground to space and between spacecraft in a standardized, unambiguous way. We consider how existing communication protocols and transfer formats could be employed to achieve this, and suggest where modifications may be necessary.

2.1. Introduction

Standardization context

Standardization is of immense importance in the aerospace domain, and particularly so for spaceflight missions, where failure is punished swiftly and harshly by the uncompromising environment of space. Clear standards could help in every part of the development pipeline, for instance: to ensure compliance and interoperability during the design phase, to establish rigorous and clear expectations for performance, to permit later integration of independently developed components, and to allow correct performance during deployment and systematic troubleshooting in case errors do occur.

Beyond ensuring the effective deployment of existing technologies, standardization is also a crucial tool to assist in the safe exploitation of new technologies, such as Artificial Intelligence and Machine Learning in the spaceflight domain. However, the standardization of machine learning presents a particular challenge, as the state of the art in the domain is progressing at such a swift pace that it is difficult to identify and formalize enduring characteristics of the technology. Compared to traditional systems, research on machine learning is evolving rapidly, requiring standards that are both flexible and robust. At the same time, these new developments in machine learning are adopted at a much more accelerated rate than technological changes have historically been embraced by the industry, further increasing the urgency of standardization development.

Use case motivation

In recent years, there has been a progressive paradigm shift in the design of space missions, moving from the traditional monolithic scenario of a single custom-designed satellite towards multi-satellite missions. Initially, this shift led to the development of multi-satellite constellations, and now mission designs are evolving beyond the traditional design towards swarms of nanosatellites. In a swarm, many satellites in formation collaborate to jointly achieve tasks beyond the capability of a single participant. This capability, coupled with the cost constraints of launching hardware into space, means that such swarms are likely to consist of many small ("nano"-) satellites with limited individual capacities.

The distributed mission design offers several advantages, including enhanced flexibility, redundancy, and cost-effectiveness granted by the possibility of re-using satellite designs and commercial off-the-shelf solutions. For example, an Earth Observation (EO) mission consisting of multiple satellites flying in formation could obtain increased coverage of the planetary surface or flexibly reallocate imaging resources to capture areas of interest in greater detail.

However, the nanosatellite era also brings new challenges, largely caused by the limited power budget and other hardware constraints of satellites. Modern satellites are capable of collecting vast amounts of data, at rates far outpacing the available communication budget; hence data cannot be streamed to a ground station for processing in real time. Yet real-time processing is crucial for many Earth Observation applications, e.g. in the detection and management of natural disasters, where the delay caused by slow processing of data could have devastating effects.

One solution to this problem is to shift the processing of data onto the satellites and automate it using machine learning models, e.g. to automatically identify the most relevant images to be forwarded to ground stations. Such approaches are currently under development in the research domain, along with many other potential applications of on-board machine learning, e.g. for communication routing, failure detection or autonomous decision-making.

For these applications to become feasible and reliable in the general sense, they should be supported by a standardization framework allowing the safe and secure development, training, fine tuning, and communication of models.

To support the utilization of machine learning models on-board spacecraft, standards about the development and deployment of such models are of crucial importance, including qualifying and quantifying the trustworthiness characteristics of systems, as discussed in a previous report [1]. Once the quality and trustworthiness of a model can be established, the next great challenge is the cross-device communication of such a model in practice. In the aerospace domain in particular, standardizing communications protocols is of crucial importance: without well-defined communications, all other aspects of a spaceflight mission are at risk. Well-defined communication standards can prevent known problems and even enhance the flexibility of a mission, allowing interactions that were not necessarily planned during the design phase of the mission, enabling compatibility between independently developed systems. With the current drive towards multi-satellite missions and the diversification of stakeholders, this aspect gains importance, as it becomes more feasible and desirable for different satellites to collaborate.

This report is dedicated to examining how standardization could support the communication aspect of such novel machine learning systems deployed on satellites, enabling a stable transfer of models from ground to satellite and between satellites and ensuring compatibility with existing space communication protocols.

2.2. Case studies

For the remainder of this paper, we focus our analysis on two broad use cases. In a first step, we consider the general problem of transmitting a single complete machine learning model from a ground station to a satellite. This corresponds to a scenario, illustrated in Figure 8, where the training of a machine learning model is carried out on the ground, using previously collected raw sensor data from the satellites or a synthetic approximation as the underlying training data. The trained model is then deployed from the ground station onto the satellite. The capability to successfully carry out this procedure – deploying and running a machine learning model on a single spacecraft – is a fundamental stepping-stone on the path towards the integration of more complex machine learning pipelines into space missions.

In the second part of our analysis, we consider one such more complex scenario building on the first: enabling on-board machine learning in distribution across multiple satellites, using the Federated Learning (FL) paradigm. Under this paradigm, each satellite trains an on-board machine learning model, periodically exchanging information about the resulting models with other satellites to enhance the learning process. This necessitates an extension of the communication standard required for the first case.

2.2.1. First case study: ground-to-satellite model transfer

In this section, we explore how a machine learning model may be deployed from the ground onto a satellite, outlining concrete approaches taken in existing missions and discussing how these relate to current standardization efforts. Given the subject of this work, our discussion will mainly focus on the network aspects of potential solutions, i.e. how a machine learning model can be encoded for transmission, and how the resulting information is transmitted.

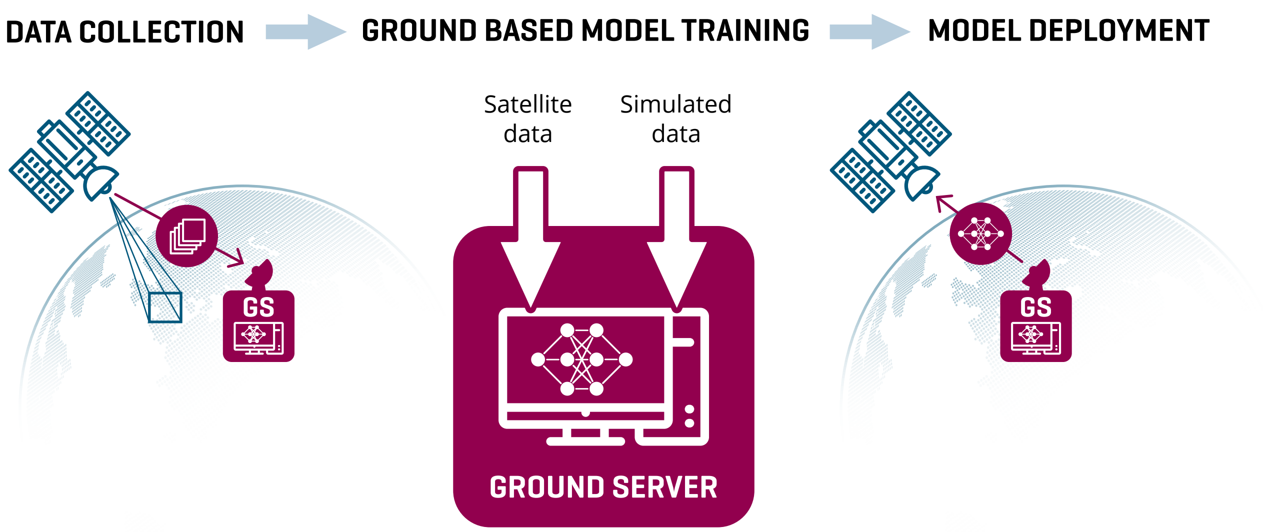

Figure 8. The scenario of the first case study. First, training data is collected from a satellite or a simulator. The centralized training data is used by a ground server to train a machine learning model (neural network). Finally, the completed model is deployed onto the satellite. In this case study, we consider how to facilitate the communication of the machine learning model from the ground to the satellite, taking place in the last step.

The deployment of machine learning models on-board spacecraft is still in its initial stages; we therefore approach our analysis from two sides. We begin by considering existing standards and methods for encoding and communicating machine learning models in the general case. Then, we discuss the solutions chosen for the Φ-sat-1 (PhiSat-1) mission[13], a recent proof-of-concept mission where a machine learning model was successfully deployed onto a satellite. Finally, we examine how a model transfer procedure could be integrated into existing space communication standards formulated by the Consultative Committee for Space Data Systems (CCSDS)[14].

Model transfer formats

Historically, machine learning models have been developed, trained, and deployed in fairly self-contained pipelines, each designed for a specific purpose and for the needs of a specific stakeholder. In practice, this has caused different programming frameworks fit for different purposes to proliferate, with no particular requirement for mutual compatibility. Not only do these frameworks encode and store models in different formats; they also translate to different behaviors upon hardware deployment. In effect, this means that models containing apparently the same abstract structures might deliver substantially different results when deployed by different frameworks.

With the recent move towards collaborative machine learning strategies and the increasing interest in deploying machine learning models across edge devices, the need for aligning different frameworks has become apparent. The absence of a unified standard introduces risks such as interoperability issues, increased development costs, and the potential for mission-critical failures. Given the entrenched differences in machine learning frameworks, the most straightforward way of facilitating alignment would be to establish a unified *model transfer format*, designed to allow models to be communicated unambiguously between frameworks.

At this time, there is no such unified model transfer format defined by an official standardization body.

Outside of the formal standardization domain, two main competing model transfer formats exist at the time of writing of this report: the Neural Network Exchange Format (NNEF)[15], and the Open Neural Network Exchange format (ONNX)[16]. Neither of the two formats can be considered an officially recognized standard, in the frame of the European Regulation (EU) No 1025/2012 on standardization[17], since both are being maintained by different industrial stakeholder consortia. Several additional tools of more limited scope exist, designed either to support specific providers, such as Intel's OpenVino tool, or to convert between specific frameworks on the application level.

For this study, we are most interested in a general solution, applicable to a variety of hardware configurations; hence we place our focus for this work on the capabilities of the NNEF and ONNX formats, in particular on how such a general solution could be integrated into the existing framework of related standardization in the relevant domains.

Closer investigation reveals that the ONNX format relies on the serialisation of the computational graph, including parameter values, into a binary format for transfer. In contrast, the Neural Network Exchange Format, maintained by the Khronos NNEF Working Group, consists of a description of the complete computational graph corresponding to the structure of the given neural network, expressed in a human-readable syntax, and the numerical parameters associated with the network. A number of common machine learning tools support the import and export of models in NNEF format. Of the two formats, NNEF appears to be preferred for potential application to the aerospace domain by domain experts, as seen e.g. in [2] due to its relatively well-documented syntax and semantics.

13 https://www.eoportal.org/satellite-missions/phisat-1

14 https://public.ccsds.org/default.aspx

15 https://www.khronos.org/nnef

16 https://onnx.ai/

17 http://data.europa.eu/eli/reg/2012/1025/oj

Existing solutions: PhiSat 1

Particularly of note for inspiration from the application domain is the PhiSat-1 (and PhiSat-2) mission[18], launched as a proof-of-concept mission by the European Space Agency (ESA). The purpose of this mission was to provide a first technology demonstration of Artificial Intelligence (Neural Networks) for Earth observation run on board a nanosatellite, including the development and deployment pipeline, hardware capability, and analysis of the results. To date, this is the only European mission to demonstrate the deployment of a pre-trained artificial Intelligence model for EO onto a satellite in space. This mission design matches our scenario; hence it is useful to consider how it was realized - both from a standardization perspective, and as a demonstration that the technology is feasible or will soon be ready for deployment, making the timing of these standardization efforts crucial.

The PhiSat-1 mission makes use of the proprietary OpenVino tool, developed specifically for Intel hardware, to facilitate the transfer of the pre-trained model onto the satellite. This choice does not generalize well in terms of standardization, as it does not cover hardware components built by other providers. From a standardization perspective, we are in contrast interested in building a more universal, hardware-agnostic standard solution.

Beyond the model transfer format, we note that the general communication with the satellite and software deployment was facilitated using the Nanosatellite MO Framework (NMF), which implements common CCSDS standards to facilitate communication. This could be of interest for further investigation into the pairing of a general model transfer format with common CCSDS standards, and the network stack could perhaps even serve as another instantiation example in addition to our case study.

Integration with CCSDS: Proposed Communications Stack

The Consultative Committee for Space Data Systems (CCSDS) defines a large number of communication protocols for use in space applications. Some CCSDS standards have been adopted as European standards within the CEN/CLC JTC 5 – *Space*[19] technical committee and as international standards within ISO/TC 20/SC 13 - *Space data and information transfer systems*[20]. Given their widespread use, it is worthwhile to consider if and how the domain-agnostic model transfer formats discussed in the previous section could be integrated into this protocol stack.

For the present use case, this appears to be quite uncomplicated: CCSDS protocols are classified following the standard Open Systems Interconnection (OSI) networking model. The transfer of a machine learning model from ground to a satellite using a model transfer format, as discussed in this chapter, could likely be encapsulated in the application layer, using a communication stack of existing CCSDS protocols. A possible instance of such a network stack is proposed in the following table, with brief citations of the related CCSDS standards to support the respective choices. This table is primarily intended to serve as an illustration; the specific choice of instantiation should be dependent on mission configuration.

OSI Layer	Example Protocol	Comment
Application	CFDP + Lossless Data Compression	"CFDP is designed to meet the needs of space missions to transfer files. It is a file transfer protocol, but it also provides services typically found in the Transport Layer, that is, complete, in-order, and without duplicate data delivery" The Lossless Data Compression standard "guarantees full reconstruction of the original data without incurring any distortion in the process. It is intended to be used together with the Space Packet Protocol or CFDP."
Transport	SCPS-TP	
Network	Bundle Protocol (BP)	

18 https://www.eoportal.org/satellite-missions/phisat-1

19 https://standards.cencenelec.eu/dyn/www/f?p=205:7:0::::FSP_ORG_ID:887985&cs=17D471F6F920904967AFC18C2BDA2F89F

20 https://www.iso.org/committee/46612.html

Data Link	Unified Space Link Protocol (USLP)	USLP has "a function for retransmitting lost or corrupted data to ensure delivery of data in sequence without gaps or duplication over a space link." CCSDS 130.0-G-4 - *OVERVIEW OF SPACE COMMUNICATIONS PROTOCOLS*[21]
Physical	CCSDS Recommended Standard for Radio Frequency and Modulation Systems	CCSDS 401.0-B-32 - *RADIO FREQUENCY AND MODULATION SYSTEMS – PART 1 - EARTH STATIONS AND SPACECRAFT*[22]

In this first case study, we have identified a standardization gap covering communication protocols for the transmission of machine learning models between ground and spacecraft, and spacecraft-to-spacecraft. We discussed existing standardization and industry standards addressing the transfer of models in other domain contexts, and how these existing formats could be encapsulated in existing space communication standards. In the next section, we will discuss a more complex use case building on the scenario discussed here.

2.2.2. Second case study: Federated Learning across satellites

In this section, we discuss the extension of communication protocols required to enable collaborative on-board machine learning across satellites. The small size and limited resources of modern nanosatellites requires the collaboration of multiple satellites to enhance on-board computing capabilities, allowing resources and information to be shared across spacecraft. We focus one such distributed machine learning paradigm, known as Federated Learning (FL), that could be particularly well-suited to this use case. Federated Learning schemes allow participants to exchange locally gathered information at a comparably low communication cost, permitting the exploitation of distributed information while respecting resource constraints.

Background: Federated Learning

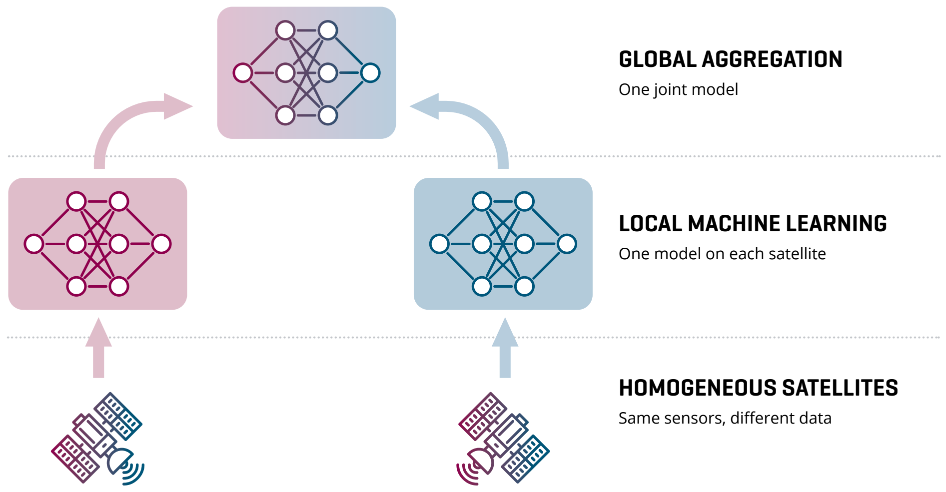

Figure 9. The classical Federated Learning paradigm deployed on satellites. Each satellite trains a local model on the data it has gathered; a server periodically aggregates the local models of multiple satellites into a single global model. Depending on the use case, the server role can be performed by a ground station or a satellite.

21 https://public.ccsds.org/Pubs/130x0g4e1.pdf

22 https://public.ccsds.org/Pubs/401x0b32.pdf

Under the classical Federated Learning strategy, illustrated in Figure 9, each distributed participant ("client", corresponding to a satellite in our scenario) trains a local machine learning model, using only its locally available data. All clients periodically share the current version of their local model with a central server – in practice, this generally corresponds to sharing the updated scalar weights assigned to each node of a neural network. The server collects the local models of multiple or all participants, aggregating these into a single global model. This global model is distributed back to the local clients to continue the learning process. Following a number of such collaborative learning steps, the global model implicitly incorporates information learned across many different clients. Research has shown that, under favorable conditions, this distributed learning scheme has the potential to ultimately deliver a global model performance matching that of a single model trained on the combined dataset of all clients [3]. Applied to our use case, this implies that multiple satellites training a model together in distribution could be capable of achieving a model of the same quality as one trained on a central Earth-bound server, but without requiring the costly step of collecting the underlying data.

However, in real-world applications, including in the space domain, Federated Learning encounters additional challenges compared to the ideal setting. Most prominently, these include imbalances caused by differences between participants, e.g. in terms of data distribution, hardware capability, participant preferences, and communication restrictions, e.g. intermittent connectivity caused by the orbital periods of satellites. Research aimed at mitigating these challenges remains ongoing, but has already led to the development of many different variants of Federated Learning algorithms, each tailored to different requirements [4].

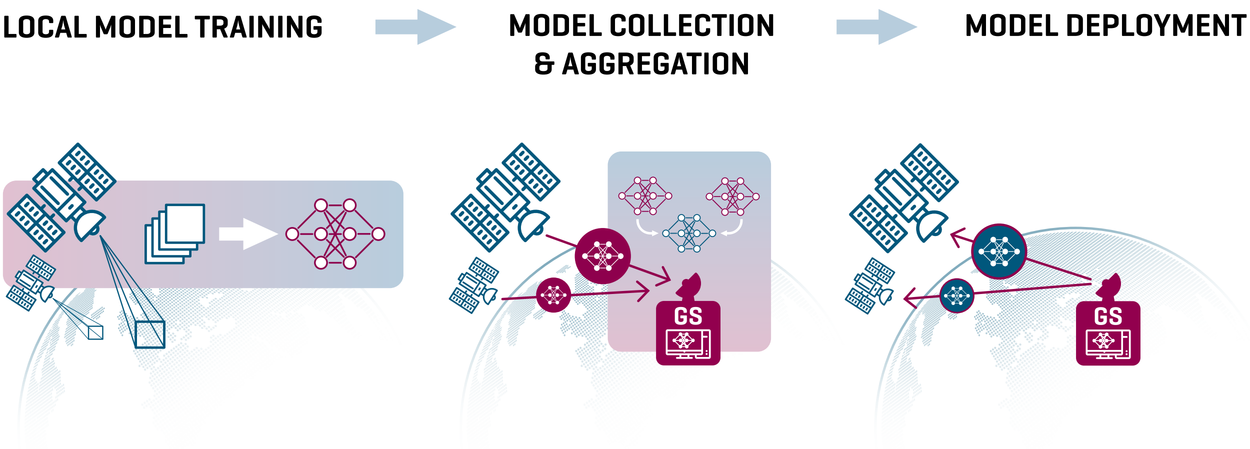

Figure 10. An example variant of the scenario of the second case study. Each participating satellite trains a local machine learning model on-board, using only its own gathered data for training (left image). Local models are collected by a central server, where the contributions of multiple satellites are aggregated into a single global model (centre). Following the server-side aggregation, the resulting global model is distributed back to all participating satellites (right).

Federated Learning Protocol

Any standardized deployment of Federated Learning, therefore, would need to begin with a clear and unambiguous communication of the specific characteristics of the algorithm being deployed. Model exchange, which can exploit one of the formats discussed in previous section, is only one of the parameters of this proposed Federated Learning protocol. We suggest the development of a dedicated Federated Learning protocol to set up and facilitate any ad-hoc collaboration between multiple satellites under the Federated Learning paradigm. No such protocol currently exists. In this section, we provide a first overview of the challenges to be solved by such a communication protocol by suggesting a list of characteristics that it should address. This list is intended to serve as a starting point; it is by no means exhaustive. Our proposals for parameters that should be included in this protocol are split into two subsets: the set of client parameters, which define the behavior of participants in the client role of the federated learning system, and the set of server parameters, which specify the behavior of the federated server.

Name of client parameter	Description	Comments
System topology	Different federated algorithms use different aggregation control schemes, e.g. fixed star topology (one central server), fully decentralised (no server at all), ...	
Server identity	Identity/address of the server, if it exists.	Needs to be fixed for clients if (1) clients are required to contact server, or (2) for security, to allow clients to verify server if contacted.
Local submission trigger	Defines how model submission is triggered on the local client: after a given number of steps, by reaching a certain training loss, upon contact by the server etc.	
Local aggregation behaviour	Defines client behaviour after local model has been submitted, but before external model update has been received. Behaviours could include pausing training until model update is received, or continuing training.	
Local integration of global update	Defines how a global model update is processed on the local client, e.g. replacing the local model, partial update, etc.	
Local submission format	Defines the type and format of local updates to be submitted by the client. E.g. weights of full model, gradient of partial model, etc.	
Initial model	Defines the architecture (or fully initialised version) of the model to be trained on the client.	Can be reduced to a set of constraints if the FL scheme does not require homogeneous client models.
Failure handling	Defines how to detect and handle different failures locally, e.g. time-out on global updates if server fails, procedure for handling server changes or local failures.	Requires detailed specification of additional fields.

Name of server parameter	Description	Comments
Global aggregation behaviour	Aggregation strategy to be used on the server.	
Model collection mechanism	Defines how models are collected from the local clients, e.g. through active collection by the server or proactive client submission.	
Initial (global) model	Defines the expected architecture of the global model.	Can be transmitted using model transfer format discussed in previous section.
Failure handling	Defines how to detect and handle different failures locally and in the system, e.g. time-out on local updates if client fails, procedure for handling server changes or local failures.	Requires detailed specification of additional fields.

Communication of Model Updates

Once the initial parameters of the federated learning scheme have been established, models must be transmitted periodically to and from the clients during the training process. In this section, we consider how these model transfer messages might be realized.

The simplest solution in terms of standardization effort would be to transmit the complete model for each update, using a model transfer format as established in the previous section. However, this would involve a significant needless expenditure of energy, as most common federated learning schemes do not modify the underlying architecture of the machine learning model during the learning process. If the underlying model structure remains fixed, the format of model updates could then be reduced to updating only those aspects of the model that do change, i.e. the scalar weights assigned to the nodes of a neural network. As the reduction of communication cost is a crucial challenge in the design of energy-efficient space missions, this possibility bears further consideration.

One of the two existing model transfer formats considered in this paper, the NNEF format, appears to be suited to this strategy with little modification required. In this format, model information is encoded in human-understandable semantics, with information about the model architecture and the parameter weights assigned to this model stored separately. Retrieving and transferring only the parameter-weight section of the encoded format, and integrating this partial update into the model on the receiving end, should present little additional difficulty.

On the other hand, implementing the same strategy with the ONNX transfer format appears to be much more complex, as this format encodes all model information in a single binary file, with no obvious way of isolating the model parameters. Doing so would likely require more significant modifications to the underlying encoding of the format than appears practical.

In conclusion, this section identified a need for the development of a dedicated communication protocol to facilitate a future use case where multiple satellites cooperate in training a machine learning model using a Federated Learning strategy. Suggestions for tackling this standardization gap in an effective manner were proposed, including a list of characteristics to consider in defining such a communication protocol, and how to leverage existing related standards to optimize the size of communication messages transmitted during the learning process.

2.3. Conclusion

In this work, we have considered the current state of standardization for the communication of machine learning models in a space context. We have analyzed two case study scenarios, one on the transfer of a trained machine learning model from a ground station onto a satellite - close to the current technological state of nanosatellite missions, such as the PhiSat-1 mission - and one representing a more ambitious near-future use case where multiple nanosatellites collaborate to perform on-board machine learning.

For the first case study, we observe that no current dedicated standard exists, but that some existing standards for transferring machine learning models, developed in the general machine learning community, could likely be combined with comparatively little additional effort to fit with established CCSDS communication standards. We have proposed an example communication stack and discussed how one such existing model transfer format, NNEF, could be encapsulated to permit the transfer of models. The NNEF transfer format is currently developed and maintained by the Khronos Group, an industry consortium; yet the current standard appears relatively well-defined and –maintained, and could perhaps be developed further in collaboration with official standardization bodies.

For the second scenario, we have identified an additional standardization gap to be overcome: the need for a well-defined protocol allowing ground- and space-based participants to negotiate their participation in a federated learning system and establish the joint parameters of this collaboration. Our analysis has yielded a broad overview of items of interest that would need to be included in the definition of such a protocol to cover a reasonable range of existing FL approaches. Finally, we have discussed how our suggested solution for the first use case could be integrated into this second scenario, and how messages defined in compliance with the NNEF model transfer format and passed during the main Federated Learning phase could be reduced effectively to conserve communication resources.

Ultimately, we have identified an achievable path towards the establishment of communication standards enabling the transfer of machine learning models in a space context in compliance with existing standardization, and demonstrated the feasibility of extending such standardization further to encourage the implementation of innovative distributed learning solutions.

References

[1] ILNAS: Trustworthiness in ICT, Aerospace, and Construction applications: scientific research and Technical Standardization (2023). ISBN 978-99987-869-7-4. Available: https://portail-qualite.public.lu/dam-assets/publications/normalisation/2023/white-paper-trustworthiness-ict-aerospace-construction-october-2023.pdf

[2] Gauffriau, A., De Albuquerque Silva, I., Pagetti, C. : Formal description of ML models for unambiguous implementation. 12th European Congress on Embedded Real Time Software and Systems (ERTS 2024). https://doi.org/10.48550/arXiv.2307.12713

[3] McMahan, B., Moore, E., Ramage, D., Hampson, S. & Arcas, B.A.y. : Communication-Efficient Learning of Deep Networks from Decentralized Data. Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, in Proceedings of Machine Learning Research 54:1273-1282 (2017). https://doi.org/10.48550/arXiv.1602.05629. Available from https://proceedings.mlr.press/v54/mcmahan17a.html.

[4] Ye, M., Fang X., Du, B., Yuen, P. C., and Tao D.: Heterogeneous Federated Learning: State-of-the-art and Research Challenges. ACM Comput. Surv. 56, 3, Article 79 (2024), 44 pages. https://doi.org/10.1145/3625558

3. Demonstrating the integration of a construction sector standard in a constraint programming model

3. Demonstrating the integration of a construction sector standard in a constraint programming model

Abstract

This report explores the integration of established technical standards into Building Information Modeling (BIM) to enhance its functionalities. Focusing on the thermal transmittance of windows, in this work we have implemented a constraint programming model using an international technical standard as input to optimize window design, balancing energy efficiency and cost-effectiveness, with a view towards helping buildings meet criteria for sustainability. The report highlights the significance of thermal transmittance in construction, the role of technical standards such as ISO 10077-1:2017 - Thermal performance of windows, doors and shutters - Calculation of thermal transmittance Part 1: General, and the application of artificial intelligence (AI) and hyperparameter optimization to streamline calculations. Results demonstrate the model's effectiveness in achieving optimal thermal performance specifications, underscoring the potential for AI-driven solutions to advance sustainable building practices using technical standards as a baseline.

3.1. Introduction

This work follows in the footsteps of the investigations described in the 2023 white paper "Trustworthiness in ICT, Aerospace, and Construction Applications - Scientific Research and Technical Standardization – October 2023" [1] on the topic of Building Information Modelling (BIM). One of the directions that was proposed (see Table 6 of [1]) to be explored in terms of linking BIM research to standardization is how to consider established standards in the construction sector as input to a given BIM model. The idea is to show how technical standards' content can be integrated into modern tools with the potential to enhance BIM functionalities and to foster the implementation of BIM on the market.

An example of how this can be achieved is detailed in the present update. Specifically, a standard on the thermal transmittance of windows is used as a baseline input to define a constraint programming model aimed at optimizing window design in a construction project as a function also of cost.

In this work, first thermal transmittance and its importance to the construction sector are introduced in the current context of heightened awareness on sustainability and energy efficiency, and the role that technical standards play is highlighted, with one standard of great relevance being showcased.

Then, how artificial intelligence impacts and enhances the information technology tool set of many domains – in particular, that of construction – is described. Additional focus is put on the usage of constraint programming and hyperparameter optimization. In the case of the former, constraint programming is found to be a natural fit to model the kinds of multi-variable equations that are used to calculate quantities such as the thermal transmittance of a window. In the case of the latter, hyperparameter optimization finds its use in greatly enhancing the fine-tuning of an AI model for it to efficiently and robustly find solutions to constraint programming problems.

Afterwards, a description of an instantiation of the standard in question through these methods is presented, along with a sample of results. To further demonstrate the flexibility of being able to add more constraints, an additional variable taking in consideration the cost of the building materials is integrated to the model.

Finally, the work is concluded with some prospects on further standards-integrating research.

3.2. Thermal transmittance, construction, and standards

3.2.1. Thermal transmittance and construction

Thermal transmittance in construction, particularly in the context of windows, doors, and shutters, is critical in building design and energy efficiency. Thermal transmittance, also known as the U-value, measures the rate of heat transfer through a building element over a given area, under standard conditions. Lower U-values indicate better insulation properties, not only contributing to reduced heating or cooling costs but also playing a significant role in improving the building's overall sustainability [2].

When designing a building's envelope, it is crucial to consider each component, as every layer has unique properties that significantly impact the building's overall thermal performance. The U-value, which is calculated by dividing 1 m² of the envelope by the temperature difference between its sides, measures this performance. Expressed in W/m²·K (Watt per Square Meter per Kelvin), the U-value indicates the level of thermal insulation based on the amount of energy passing through. A low U-value means the surface is well-insulated, while a high U-value indicates poor thermal insulation. This value is influenced by the thermal resistance of each layer, which depends on the thickness and thermal conductivity of the materials used[23].

The importance of thermal transmittance in construction cannot be overstated. It directly impacts the energy performance of buildings, influencing both the comfort of occupants and the operational costs associated with heating and cooling. Energy performance is the overall effectiveness of a building in using energy to provide comfort and functionality. In regions with extreme climates, the role of thermal transmittance becomes even more critical, as it can significantly affect the energy demand and sustainability of buildings. By focusing on improving the U-values of windows (shown in Figure 1), architects and engineers can design buildings that are not only more energy-efficient but also more resilient to climate change.

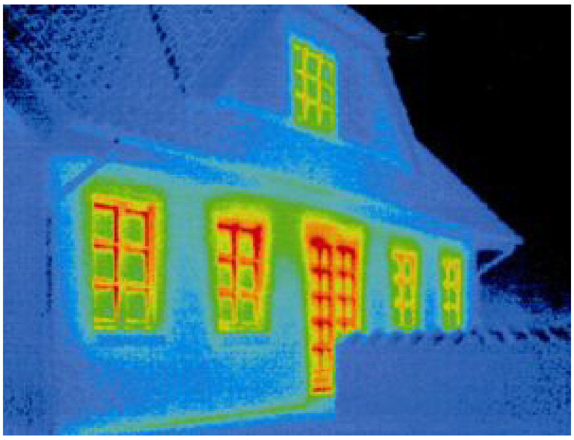

Figure 11: This thermographic image reveals that the highest heat loss (red) takes place through the window frames. It is estimated that more than one third of the energy used to heat buildings is lost through windows and doors[24].

According to the EU taxonomy[25] and Regulation (EU) 2020/852[26], environmentally sustainable economic activities are those that contribute to climate change mitigation, including energy efficiency. Therefore, improved thermal

23 https://www.archdaily.com/898843/how-to-calculate-the-thermal-transmittance-u-value-in-the-envelope-of-a-building

24 https://miwindows.com/

25 https://finance.ec.europa.eu/sustainable-finance/tools-and-standards/eu-taxonomy-sustainable-activities_en

26 https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A32020R0852&qid=1724405839677

transmittance aligns with these sustainable activities, as it supports making buildings more energy-efficient and environmentally friendly. This broader perspective highlights the importance of thermal transmittance beyond just cost savings, emphasizing its role in promoting sustainability in the construction industry [3]. According to UNEP's "2022 Global Status Report for Buildings and Construction", the construction industry accounts for approximately 38% of global energy-related CO_2 emissions[27].

In addition to the EU taxonomy, various national and international regulations emphasize the need for mandating zero-emission buildings for all new constructions and major renovations. The European Commission introduced a directive (the Energy Performance of Buildings Directive, or EPBD - 2010/31/EU) aimed at significantly reducing the carbon footprint of the building sector[28]. By setting strict energy efficiency rules and promoting the use of renewable energy sources, the directive seeks to ensure that buildings are not only energy-efficient but also capable of producing as much energy as they consume. This move is part of the broader European Green Deal, which aims to make Europe the first climate-neutral continent by 2050[29].

In achieving these ambitious targets, a critical focus is placed on the thermal performance of building elements, as it directly impacts a building's energy efficiency and overall carbon footprint.

3.2.2. The role of technical standards

In general, tackling global-scale problems such as sustainability requires far-reaching agreement on how to describe the problems to be solved, and through which methodologies to solve them. Indeed, having one effective solution agreed upon by many allows for all to be working more effectively towards a single common goal, thus avoid duplication of efforts and fostering overall interoperability.

The process by which technical standards are developed aims exactly at 1) reaching this kind of agreement across a broad range of stakeholders in each domain and 2) producing the kinds of reference documents that describe the obtained solutions. Hence, relying on technical standards as inputs is a particularly useful exercise.

In the specific case of thermal transmittance of windows, there exists an international standard that provides a base formula and agreed methods to select appropriately the parameters to be used depending on the desired configuration. This can be a function of location, weather conditions, materials used, and many other variables. The standard in question is ISO 10077-1:2017 - *Thermal performance of windows, doors and shutters – calculation of thermal transmittance – Part 1: General*[30], created and maintained by the international technical standardization committee ISO/TC 163/SC 2 - *Calculation methods*[31] (ISO/TC 163 is *Thermal performance and energy use in the built environment*[32]).

Calculating thermal transmittance is very complex due to the various materials and components involved in a building element, and ISO 10077-1:2017 provides detailed methodologies. However, these methods often require extensive data and can be computationally intensive. Thermal transmittance and the calculation of U-value are crucial for determining a building's energy performance, which is necessary for obtaining an energy certificate – a document required for various types of construction projects. Each country has its own requirements for energy performance. For example, in Luxembourg, the regulations are outlined in the *Règlement grand-ducal du 9 juin 2021 concernant la performance énergétique des bâtiments*[33].

To address these computational challenges and enhance the efficiency of calculations, advanced technologies such as Artificial Intelligence (AI) are increasingly being employed.

27 https://www.unep.org/news-and-stories/press-release/co2-emissions-buildings-and-construction-hit-new-high-leaving-sector

28 https://www.europarl.europa.eu/factsheets/en/sheet/69/efficacite-energetique (the revised EPBD (EU/2024/1075) entered into force in 2024)

29 https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficient-buildings/energy-performance-buildings-directive_en

30 https://ilnas.services-publics.lu/ecnor/displayStandard.action?id=198682

31 https://www.iso.org/committee/53512.html

32 https://www.iso.org/committee/53476.html

33 https://legilux.public.lu/eli/etat/leg/rgd/2021/06/09/a439/jo#intituleAct

3.3. Leveraging artificial intelligence

Artificial Intelligence (AI) refers to the simulation of human intelligence in machines that are designed to think and act like humans. These machines are programmed to perform tasks such as learning, reasoning, problem-solving, perception, and language understanding. AI can be categorized into two types: narrow AI, which is designed for a specific task, and general AI, which aims to perform any intellectual task that a human can [4].

AI technologies are primarily based on machine learning, where systems learn from data, identify patterns, and make decisions with minimal human intervention. Key techniques in AI include supervised learning, unsupervised learning, and reinforcement learning. AI is widely used in various fields, including healthcare, finance, transportation, entertainment, and especially construction, making it a crucial part of modern technology.

AI offers innovative solutions to streamline and optimize the analysis of thermal transmittance, enabling more accurate and faster predictions. The integration of AI in the analysis of thermal transmittance represents a significant technological advancement. AI algorithms can process vast amounts of data from different building components, identifying patterns and correlations that might be missed by traditional methods. This capability not only enhances the accuracy of thermal transmittance calculations but also allows for real-time adjustments during the design and construction phases. As a result, AI-driven solutions can lead to more efficient and sustainable building practices.

The complexity of the thermal transmittance calculation motivates the need for a more streamlined approach. By modelling it as a Constraint Optimization Problem (COP), we can leverage advanced solvers to automate and optimize the process. This not only simplifies the calculation but also ensures accuracy and efficiency, aligning with modern computational capabilities and standards compliance.

As AI continues to evolve, it often intersects with other computational methods to enhance its capabilities. One such method is Constraint Programming (CP), which offers a powerful approach to solving complex combinatorial problems that AI systems frequently encounter.

3.3.1. Encoding complex equations with constraint programming

CP is a method used in computer science to solve combinatorial problems that involve many variables and constraints. CP focuses on stating the conditions that any solution to the problem must satisfy, and then using algorithms to find solutions that meet these conditions [5].

A constraint satisfaction problem (CSP) consists of three main components: a set of variables, a domain of values for each variable, and a set of constraints that specify allowable combinations of values. CP is particularly powerful for solving scheduling, planning, and resource allocation problems [6]. CP employs techniques such as backtracking, constraint propagation, and heuristics to efficiently navigate the search space and find optimal or near-optimal solutions. Integrating CP with AI can greatly improve problem-solving efficiency and solution quality. However, achieving optimal performance often requires fine-tuning various parameters of the models and algorithms used.

By modeling the thermal transmittance calculation as a Constraint Optimization Problem (COP), we can leverage advanced solvers to automate and optimize the process.

3.3.2. Using hyperparameter optimization to effectively tackle constraint optimization

Hyperparameter optimization is the process of tuning the hyperparameters of a machine learning model to improve its performance. Hyperparameters are the configuration settings used to structure and optimize a model, such as learning rate, number of layers in a neural network, and regularization parameters [7].

Unlike model parameters, which are learned during the training process, hyperparameters are set before training begins and significantly influence the model's behavior and performance. The process involves selecting a set of hyperparameters, training the model, evaluating its performance, and iterating this process to find the optimal settings.

Common techniques for hyperparameter optimization include grid search, random search, and more advanced methods like Bayesian optimization and genetic algorithms [8].

3.4. Model instantiation and results

3.4.1. The thermal transmittance equation and its variables

The ISO 10077-1:2017 standard outlines methods to calculate the thermal transmittance of windows, doors, and shutters. The standard provides equations to account for various components such as the frame, glazing, and edges, ensuring a comprehensive assessment of thermal performance.

The thermal transmittance of a single window (U_w) is calculated using the formula:

$$U_W = \frac{\sum A_g * U_g + \sum A_f * U_f + \sum l_g * \Psi_g + \sum l_{gb} * \Psi_{gb}}{A_f + A_g}$$

In this equation:

- $\mathbf{A_g}$ represents the glazed area,
- $\mathbf{U_g}$ is the thermal transmittance of the glazing (single or multiple),
- $\mathbf{A_f}$ is the sum of the frame area,
- $\mathbf{U_f}$ is the thermal transmittance of the frame,
- $\mathbf{l_g}$ is the length of the glazing edge,
- $\mathbf{\Psi_g}$ is the linear thermal transmittance of the glazing edge (due to combined effect of glazing, spacer and frame. In the case of single glazing, the linear thermal transmittance of the glazing (Ψg) is taken as zero because any correction is negligible),
- $\mathbf{\Psi_{gb}}$ is the linear thermal transmittance of the glazing bar, and
- $\mathbf{l_{gb}}$ is the length of the glazing bar.

This equation ensures that the window's overall thermal transmittance is a weighted average of the thermal transmittance of the glazing and frame, along with the edges' and bars' linear thermal transmittance.

Our COP model captures these variables and constraints to compute $\mathbf{U_w}$. The model uses an XML format compatible with XCSP3[34] standards, ensuring clear definition and easy integration with solvers. Each variable is defined with realistic ranges, and the constraint captures the core equation for $\mathbf{U_w}$. The objective is to minimize $\mathbf{U_w}$, reflecting the goal of achieving optimal insulation.

Here, the variables are presented as part of the implemented model, along with proposed ranges for each variable:

Variable	Metrics	Explanation	Range
A_g	Area(m²)	Glazed area	1...5 (m²)
A_f	Area(m²)	Frame area	0.5...2 (m²)
U_g	Thermal trans-mittance (W/m²K)	Thermal transmittance of glazing	1...3 (W/m²K)
U_f	Thermal trans-mittance (W/m²K)	Thermal transmittance of frame	2...4 (W/m²K)
psi_g	Linear thermal transmittance(W/mK)	Linear thermal transmittance of glazing edge	0.03...0.1 (W/mK)
l_g	Length(m)	Length of glazing edge	5...10 (m)
psi_{gb}	Linear thermal transmittance(W/mK)	Linear thermal transmittance of glazing bar	0...0.01 (W/mK)
l_{gb}	Length(m)	Length of glazing bar	0...1 (m)
U_w	Thermal trans-mittance(W/m²K)	Thermal transmittance of the window	0...3 (W/m²K)
M_f	Material	Material of the frame	0...3 (0: Wood, 1: Aluminum, 2: PVC, 3: Composite)
M_g	Material	Material of the glazing	0...3 (0: Single Glazing, 1: Double Glazing, 2: Triple Glazing, 3: Low-E Glazing)
C_f	Cost (€)	Cost of the frame material	10...100 (€)
C_g	Cost (€)	Cost of the glazing material	20...100 (€)

3.4.2. Model implementation

Implementing the model involved several technical steps, starting with the selection of tools and frameworks. We used the XCSP3 format for defining the COP model, ensuring compatibility with solvers like CHOCO [9]. The model was represented in an XML file, where variables, constraints, and objectives were explicitly defined.

The variables included the areas of glazing and frame (note that, as seen in Figure 1, modeling the frame is particularly important, as this is where a great deal of heat passes), their respective thermal transmittance values, and linear transmittance coefficients for edges and bars. Each variable was given a specific range based on typical values found in building elements.

In addition to thermal transmittance, cost is a critical factor in construction decisions. By including cost variables, we aim to balance energy efficiency with financial feasibility, ensuring that the building is both sustainable and cost-effective. The model allows for the calculation of both U-value and cost, or just one of them, depending on the specific requirements of the project.

34 https://xcsp.org/

Extra variables that are defined, are as follows:

- **U_w** (Thermal transmittance of the window), With a range of 0 to 3 W/m²K.
- **M_f** (Material of the frame), Represented by integers (0: Wood, 1: Aluminum, 2: PVC, 3: Composite).
- **M_g** (Material of the glazing), Represented by integers (0: Single Glazing, 1: Double Glazing, 2: Triple Glazing, 3: Low-E Glazing).
- **C_f** (Cost of the frame material), With a range of 10 to 100 €.
- **C_g** (Cost of the glazing material), With a range of 20 to 100 €.

The core of the implementation is defining the constraint equation. This was done using an intension constraint in the XML file, which accurately captured the relationship between the variables to compute **U_w** (In this model, some of the variables are displayed with their values scaled by a factor of 100 for the calculation purposes).

```
```xml
<constraints>
  eq(U_w[0], div(add(add(mul(A_g[0], U_g[0]), mul(A_f[0], U_f[0])), add(mul(psi_g[0], l_g[0]), mul(psi_gb[0], l_gb[0]))), add(A_f[0], A_g[0]))) <!-- Thermal transmittance of the window (U_w) -->
  <intension>
    or(
      and(eq(M_f[0], 0), eq(U_f[0], 300), eq(C_f[0], 5000)),
      and(eq(M_f[0], 1), eq(U_f[0], 400), eq(C_f[0], 7000)),
      and(eq(M_f[0], 2), eq(U_f[0], 200), eq(C_f[0], 4000)),
      and(eq(M_f[0], 3), eq(U_f[0], 100), eq(C_f[0], 9000))
    ) <!-- Thermal transmittance and cost of the frame (U_f, C_f) -->
</intension>
  <intension>
    or(
      and(eq(M_g[0], 0), eq(U_g[0], 500), eq(C_g[0], 3000)),
      and(eq(M_g[0], 1), eq(U_g[0], 300), eq(C_g[0], 5000)),
      and(eq(M_g[0], 2), eq(U_g[0], 200), eq(C_g[0], 7000)),
      and(eq(M_g[0], 3), eq(U_g[0], 100), eq(C_g[0], 9000))
    ) <!-- Thermal transmittance and cost of the glazing (U_g, C_g) -->
 </intension>
    <intension>
       and(ge(A_f[0], div(A_g[0], 10)), le(A_f[0], div(A_g[0], 3))
    ) <!-- Ensure the frame area is between 10% and 33.3%of the glazed area -->
  </intension>
</constraints>
```

To provide a clearer understanding, the provided XML code snippet defines the constraints for calculating the thermal transmittance of the window (**U_w**) which combines the thermal transmittance of the glazing (**U_g**) and the frame (**U_f**), weighted by their respective areas (**A_g** and **A_f**). It also includes the linear transmittance coefficients for the edges (**psi_g**) and bars (**psi_{gb}**), multiplied by their lengths (**l_g** and **l_{gb}**), The result is then divided by the total area of the frame and glazing.

The XML also includes equations that account for the areas of glazing and frame, their respective thermal transmittance values, and linear transmittance coefficients for edges and bars. A further additional constraint was added to ensure that the frame area is between 10% and 33.3% of the glazed area. Additionally, it specifies
```

the thermal transmittance and cost for different materials used in the frame and glazing, allowing the model to balance energy efficiency and cost.

The thermal transmittance and cost of the frame are determined based on the material of the frame ($\mathbf{M_f}$):

If $\mathbf{M_f}$ = 0 (Wood), then $\mathbf{U_f}$ = 3 W/m²K and $\mathbf{C_f}$ = 50 €.
If $\mathbf{M_f}$ = 1 (Aluminum), then $\mathbf{U_f}$ = 4 W/m²K and $\mathbf{C_f}$ = 70 €.
If $\mathbf{M_f}$ = 2 (PVC), then $\mathbf{U_f}$ = 2 W/m²K and $\mathbf{C_f}$ = 40 €.
If $\mathbf{M_f}$ = 3 (Composite), then $\mathbf{U_f}$ = 1 W/m²K and $\mathbf{C_f}$ = 90 €.

Also, the thermal transmittance and cost of the glazing are determined based on the material of the glazing ($\mathbf{M_g}$):

If $\mathbf{M_g}$ = 0 (Single Glazing), then $\mathbf{U_g}$ = 5 W/m²K and $\mathbf{C_g}$ = 30 €.
If $\mathbf{M_g}$ = 1 (Double Glazing), then $\mathbf{U_g}$ = 3 W/m²K and $\mathbf{C_g}$ = 50 €.
If $\mathbf{M_g}$ = 2 (Triple Glazing), then $\mathbf{U_g}$ = 2 W/m²K and $\mathbf{C_g}$ = 70 €.
If $\mathbf{M_g}$ = 3 (Low-E Glazing), then $\mathbf{U_g}$ = 1 W/m²K and $\mathbf{C_g}$ = 90 €.

These are some common values that are mostly assigned to such variables. It can include any range that the user needs[35,36].

The objective of the model is to minimize the thermal transmittance of the window ($\mathbf{U_w}$) as well as the cost of it ($\mathbf{C_f}$ + $\mathbf{C_g}$). This was set as a minimize expression in the objectives section of the XML file.

By defining these variables, constraints, and objectives, the model ensures that the optimal combination of materials and dimensions is selected to achieve the lowest possible thermal transmittance for the window, thereby improving energy efficiency.

During implementation, challenges such as ensuring correct syntax and compatibility with the solver were addressed. Testing was conducted to validate the model, using sample data to verify that the solver could correctly parse and solve the problem. The final XML model was successfully parsed and solved by Choco solver, demonstrating the viability of our approach.

3.4.3. Results

The results from solving the COP model provided insights into the thermal transmittance of the window. By inputting specific values for the variables, the solver calculated the optimal $\mathbf{U_w}$, reflecting the window's thermal performance.

In the conducted research, the XCSP3 format was utilized, which is a standard for representing constraint satisfaction and optimization problems. The problem was defined with several variables related to the window's construction. These variables included the glazed area, frame area, thermal transmittance of glazing and frame, linear thermal transmittance of glazing edge and bar, length of glazing edge and bar, thermal transmittance of the window, material of the frame and glazing, and cost of the frame and glazing materials.

The constraints of the problem were mathematical equations that the variables needed to satisfy. For instance, the thermal transmittance of the window ($\mathbf{U_w}$) was calculated as a function of the other variables.

The objectives of the problem were to minimize the thermal transmittance of the window ($\mathbf{U_w}$) and the total cost ($\mathbf{C_f}$ + $\mathbf{C_g}$).

35 http://www.passiv.de/komponentendatenbank/files/pdf/uebergang/zd_rehau_rehau-geneo-phz_en.pdf
36 https://en.wikipedia.org/wiki/Thermal_transmittance

The Choco solver (version 4.10.14) was used to find an optimal solution for minimizing the thermal transmittance of the window (**U_w**). The solver successfully found an optimal solution with an objective value of 111. This value is scaled by a factor of 100 for computational purposes, meaning the actual thermal transmittance of the window (**U_w**) in the optimal solution is 1.11 W/m²K.

The values of the variables in the optimal solution were as follows:

Variable	Value	Explanation
A_g	4.92	Glazed area (m²)
A_f	0.5	Frame area (m²)
U_g	1	Thermal transmittance of glazing (W/m²K)
U_f	2	Thermal transmittance of frame (W/m²K)
psi_g	0.03	Linear thermal transmittance of glazing edge (W/mK)
l_g	5	Length of glazing edge (m)
psi_{gb}	0	Linear thermal transmittance of glazing bar (W/mK)
l_{gb}	0	Length of glazing bar (m)
U_w	1.11	Thermal transmittance of the window (W/m²K)
M_f	2	Material of the frame (PVC)
M_g	3	Material of the glazing (Low-E Glazing)
C_f	40	Cost of the frame material (€)
C_g	90	Cost of the glazing material (€)

The optimal solution found by the Choco solver minimizes the thermal transmittance of the window to 1.11 W/m²K. This solution uses a combination of PVC for the frame and Low-E Glazing for the glazing, with specific values for the areas and thermal transmittance properties. The total cost for the materials is 130 € (40 € for the frame and 90 € for the glazing).

The detailed results are as follows (In this part of the results, some of the variables are displayed with their values scaled by a factor of 100 for calculation purposes):

```
c Choco 231102

o 111 0.1

s OPTIMUM FOUND

v <instantiation id='sol2' type='solution' cost='111' >
v     <list>A_g[0] A_f[0] U_g[0] U_f[0] psi_g[0] l_g[0] psi_gb[0] l_gb[0] U_w[0] M_f[0] M_g[0] C_f[0] C_g[0] </list>
v     <values>492 50 100 200 3 500 0 0 111 2 3 4000 9000 </values>
v </instantiation>

d FOUND SOLUTIONS 2
```

The solver found a solution with the value of 111 (1.11 W/m²K). The second number (0.1) represents the time taken to find this solution.

The solver reported that this was the optimal solution, meaning no better solution exists within the given constraints and variable ranges.

The solver explores multiple possible solutions within the defined constraints and variable ranges.

It suggests the best possible solution it finds, but if there are more, the users can choose different options based on their specific needs or preferences. Sometimes, the solver may return "unsatisfiable," indicating that no solution meets all the constraints within the given variable ranges especially if the variable values are set to the ranges that are not logical or compatible with the constraints.

For example, if we change the ranges of the frame area as follows:

Variable	Metrics	Explanation	Range
$\mathbf{A_f}$	Area(m^2)	Frame area	2...4 (m^2)

the solver will return "unsatisfiable" because the new range for frame area ($\mathbf{A_f}$) will not meet the constraint that the frame area should be between 10% and 33.3% of the glazed area. This highlights the importance of setting logical and compatible ranges for the variables to ensure that a feasible solution can be found.

3.5. Conclusion and further prospects

The work aimed to optimize the thermal transmittance ($\mathbf{U_w}$) of windows by leveraging the COP model and using the ISO 10077-1:2017 standard. The model's objective was to minimize $\mathbf{U_w}$ while considering various parameters such as glazed and frame areas, thermal transmittance of glazing and frame, linear thermal transmittance of glazing edges and bars, and the cost and materials of the frame and glazing.

By defining the problem in the XCSP3 format and employing the Choco solver, the model successfully found an optimal solution. The solver minimized the thermal transmittance of the window, and the total material cost for this configuration for each variable contributing to the optimal solution. The solver found a solution with a thermal transmittance of 1.11 W/m²K and a cost of 130 €.

The model's results provide significant insights into achieving energy-efficient window designs. By optimizing the combination of materials and dimensions, the research highlights the potential for substantial improvements in thermal performance, which is crucial for enhancing building energy efficiency.

The research underscores the importance of material selection and precise parameter definition in achieving energy-efficient window designs.

Future research could expand on these findings by exploring additional variables and materials, validating the model against real-world data, and applying the optimization approach to other building components. Indeed, construction domain standards cover an enormous range of topics, from physical considerations (such as thermal transmittance) including structural resistance, parameters on concrete or steel, and luminescence to managerial aspects of construction, such as scheduling and safety. It remains to explore which of these can be effectively integrated into AI constraint programming models, and at which relevant levels, to then have solvers such as the one prototyped here effectively search for optimal solutions.

It should be noted that since this work is more orientated towards utilizing standards focusing more on the engineering and physical aspects of the construction sector as input to CP-driven AI models, it could be of interest to the BIM standardization community, in particular those involved for instance in committees ISO/TC 59/SC 13 - *Organization and digitization of information about buildings and civil engineering works, including building information modelling (BIM)*[37] or CEN/TC 442 - *Building Information Modelling (BIM)*[38].

37 https://www.iso.org/committee/49180.html

38 https://standards.cencenelec.eu/dyn/www/f?p=205:7:0::::FSP_ORG_ID:1991542&cs=100E563A3950D53807585F6A443ACB202

References

[1] ILNAS : Trustworthiness in ICT, Aerospace, and Construction applications: scientific research and Technical Standardization (2023). ISBN 978-99987-869-7-4. Available: https://portail-qualite.public.lu/dam-assets/publications/normalisation/2023/white-paper-trustworthiness-ict-aerospace-construction-october-2023.pdf

[2] Bucklin, O., Müller, T., Amtsberg, F., Leistner, P., Menges, A. : Analysis of thermal Transmittance, air Permeability, and hygrothermal behavior of a solid timber building envelope. Energy and Buildings, Volume 299, 113629 (2023). https://doi.org/10.1016/j.enbuild.2023.113629.

[3] Rabani, M., Bayera Madessa, H., Nord, N. : Achieving zero-energy building performance with thermal and visual comfort enhancement through optimization of fenestration, envelope, shading device, and energy supply system. Sustainable Energy Technologies and Assessments 44, 101020 (2021). https://doi.org/10.1016/j.seta.2021.101020

[4] Russell, S., & Norvig, P. : Artificial Intelligence: A Modern Approach. Pearson (2016). ISBN: 0-13-461099-7

[5] Rossi, F., van Beek, P., & Walsh, T. : Handbook of Constraint Programming. Elsevier (2006). ISBN: 9780444527264

[6] Apt, K. R. : Principles of Constraint Programming. Cambridge University Press (2003). https://doi.org/10.1017/CBO9780511615320

[7] Feurer, M., Hutter, F. : Hyperparameter Optimization in Automated Machine Learning. Springer (2009). https://doi.org/10.1007/978-3-030-05318-5_1

[8] Bergstra, J.,Bengio, Y. : Random Search for Hyper-Parameter Optimization. Journal of Machine Learning Research, 13(Feb), 281-305 (2012).

[9] Prud'homme, C., Fages, J.-G. : Choco-solver: A Java library for constraint programming. Journal of Open-Source Software 7(78) (2011). http://dx.doi.org/10.21105/joss.04708. Available: https://github.com/chocoteam/choco-solver

October 2024 · TECHNICAL REPORT · RESEARCH-DRIVEN APPROACH TO STANDARDIZATION · ICT, CONSTRUCTION AND AEROSPACE. · Version 1.0 · © ILNAS/University of Luxembourg/ANEC GIE/SnT · ISBN 978-99987-734-3-1

ILNAS
Institut Luxembourgeois de la Normalisation, de l'Accréditation, de la Sécurité et qualité des produits et services

uni.lu
UNIVERSITÉ DU LUXEMBOURG

Southlane Tower I · 1, avenue du Swing · L-4367 Belvaux · Tel. : (+352) 24 77 43 -40 · Fax : (+352) 24 79 43 -70 · E-mail : normalisation@ilnas.etat.lu

www.portail-qualite.lu