Jan, 2026
GRINS DISCUSSION PAPER SERIES DP N° 79/2026 ISSN 3035-5567

GRINS

Development Of An Open Access Algorithm To Extract Data Relating To Solar Exposure From Cadastral Maps

DP N° 79/2026

Authors:
Tlepin Sanjar

Jan, 2026
GRINS DISCUSSION PAPER SERIES DP N° 79/2026 ISSN 3035-5567

Development Of An Open Access Algorithm To Extract Data Relating To Solar Exposure From Cadastral Maps

Tlepin Sanjar

KEYWORDS

Cadastral maps | Dwellings' Solar Exposure | Computer Vision | Visual Undestanding | Data Extraction

JEL CODE

C63

ACKNOWLEDGEMENTS

This study was funded by the European Union - NextGenerationEU, in the framework of the GRINS - Growing Resilient, INclusive and Sustainable project (GRINS PE00000018). The views and opinions expressed are solely those of the authors and do not necessarily reflect those of the European Union, nor can the European Union be held responsible for them.

CITE THIS WORK

Author(s): Tlepin Sanjar. Title: Development Of An Open Access Algorithm To Extract Data Relating To Solar Exposure From Cadastral Maps. Publication Date: 2026.

The aim of this work is to develop an automated pipeline that extracts solar direc- tional information of the buildings (i.e., the orientation of windows, rooms, walls and balconies) from cadastral scans. This work is going to support the design of energy efficiency policy in Italian residential public housing, a setting where public resources are typically scarce and tenants suffer from different vulnerabilities.
Specifically, in the GRINS project, Spoke 6, work package 2, the research line on energy poverty has developed two pilot surveys to collect information from public housing tenants in Reggio Emilia and in Foggia on the energy efficiency of their homes, their energy expenditure, and their perception of thermal comfort.
The present work will com- plement the surveys' outcomes by adding relevant information from cadastral scans. This information has to be structured, so the needed energy-environmental analyses to be developed by merging different databases, can deliver key elements to design targeted and effective actions. By applying Artificial Intelligence (AI) methods (i.e. models that solve detection and classification tasks), the model generates a dataset (CSV/Excel) suitable for analyses of sunlight exposure. The model to extract information is designed to minimize manual labeling, preserve anonymity, and enable scalability for large volumes of incoming scans.
Given the approach adopted to extract information, the present model can be usefully adapted to complement information in settings other than public residential housing, where policies to

reduce carbon emissions from housing need to be planned. From a methodological standpoint, the present model builds on state-of-the-art computer vision techniques. The YOLO (You Only Look Once), namely the object detection architecture, is widely recognized for its speed and accuracy, and has been successfully applied in domains such as autonomous driving, surveillance, and document image analysis.
Similarly, ResNet architectures represent a milestone in deep convolutional networks, providing strong performance in visual classification and orientation tasks. In document understanding, deep learning approaches such as LayoutLM and DiT have demonstrated effectiveness in parsing complex layouts, though most focus on text, tables, or diagrams rather than architectural scans. The novelty of the current model lies in both its coding integration and appli- cation. Technically, it unifies object detection (YOLO) and orientation classifica- 2 tion (ResNet18) into a single end-to-end pipeline tailored to cadastral scans, going beyond conventional document parsing by coupling layout detection with compass- based directional mapping.
From an application perspective, it offers an automated solution for transforming raw cadastral data into machine-readable datasets, enabling scalable and anonymous analysis in support of sustainable urban policies, among other potential applications. To the best of our knowledge, up to now, the leveraging with geospatial data has not been widely explored in the literature and it can thus represent a completely new field for further investigations and related applications. This combination of technical rigor and practical relevance underlines the novelty of the current model in the fields of document understanding and applied computer vision. The paper is organ- ised as follows: in Section 2 we describe existed literature that can be used to solve our task ; in Section 3, we discus what methods and why we choose them to solve the problem; in Section 4 we demonstrate results and have some discussion on them with ideas for the future; Finally, Section 5 concludes all paper.

DEVELOPMENT OF AN OPEN ACCESS ALGORITHM TO EXTRACT DATA RELATING TO SOLAR EXPOSURE FROM CADASTRAL MAPS

Tlepin Sanjar
Department of Economics and Management "Marco Fanno",
University of Padova, Italy.
sanjar.tlepin@unipd.it

1 Introduction

The aim of this work is to develop an automated pipeline that extracts solar directional information of the buildings (i.e., the orientation of windows, rooms, walls and balconies) from cadastral scans. This work is going to support the design of energy efficiency policy in Italian residential public housing, a setting where public resources are typically scarce and tenants suffer from different vulnerabilities. Specifically, in the GRINS project, Spoke 6, work package 2, the research line on energy poverty has developed two pilot surveys to collect information from public housing tenants in Reggio Emilia and in Foggia on the energy efficiency of their homes, their energy expenditure, and their perception of thermal comfort. The present work will complement the surveys' outcomes by adding relevant information from cadastral scans.

This information has to be structured, so the needed energy-environmental analyses to be developed by merging different databases, can deliver key elements to design targeted and effective actions.

By applying Artificial Intelligence (AI) methods (i.e. models that solve detection and classification tasks), the model generates a dataset (CSV/Excel) suitable for analyses of sunlight exposure. The model to extract information is designed to minimize manual labeling, preserve anonymity, and enable scalability for large volumes of incoming scans.

Given the approach adopted to extract information, the present model can be usefully adapted to complement information in settings other than public residential housing, where policies to reduce carbon emissions from housing need to be planned.

From a methodological standpoint, the present model builds on state-of-the-art computer vision techniques. The YOLO (You Only Look Once), namely the object detection architecture [1], is widely recognized for its speed and accuracy, and has been successfully applied in domains such as autonomous driving, surveillance, and document image analysis. Similarly, ResNet architectures [2] represent a milestone in deep convolutional networks, providing strong performance in visual classification and orientation tasks.
In document understanding, deep learning approaches such as LayoutLM [3] and DiT [4] have demonstrated effectiveness in parsing complex layouts, though most focus on text, tables, or diagrams rather than architectural scans.

The novelty of the current model lies in both its coding integration and application. Technically, it unifies object detection (YOLO) and orientation classifica-

tion (ResNet18) into a single end-to-end pipeline tailored to cadastral scans, going beyond conventional document parsing by coupling layout detection with compass-based directional mapping.

From an application perspective, it offers an automated solution for transforming raw cadastral data into machine-readable datasets, enabling scalable and anonymous analysis in support of sustainable urban policies, among other potential applications. To the best of our knowledge, up to now, the leveraging with geospatial data has not been widely explored in the literature and it can thus represent a completely new field for further investigations and related applications. This combination of technical rigor and practical relevance underlines the novelty of the current model in the fields of document understanding and applied computer vision.

The paper is organised as follows: in Section 2 we describe existed literature that can be used to solve our task ; in Section 3, we discus what methods and why we choose them to solve the problem; in Section 4 we demonstrate results and have some discussion on them with ideas for the future; Finally, Section 5 concludes all paper.

2 Related Lituarture

In this section we discuss our research question and the related literature, splitting it in three parts: 3.1) document understanding, 3.2) solar radiation evaluation, 3.3) applications.

2.1 Document Understanding and Visual Information

The problem we address here lies within the field of visually rich document understanding. This field is relevant not only to industrial and business applications but also to academic research. One of the most prominent approaches in this domain is the *LayoutLM* [3], a transformer-based architecture that jointly models textual, spatial, and visual information for document understanding tasks. Unlike traditional language models that rely solely on textual input, LayoutLM incorporates the two-dimensional layout of text within documents, combining token embeddings with positional coordinates.

The LayoutLM model is pre-trained on large-scale document corpora using masked language modeling and next-sentence prediction objectives, similar to BERT [5], but extended with spatial embeddings that represent text-box positions. This design enables LayoutLM to effectively capture relationships between text and layout, making it particularly suitable for tasks such as form understanding, receipt parsing,

and document classification. Empirical results show that LayoutLM outperforms text-only baselines on benchmarks such as FUNSD [6] and RVL-CDIP [7], demonstrating the importance of integrating visual and layout features for document image understanding.

While LayoutLM demonstrates strong performance in extracting and interpreting textual information from documents (e.g., forms, transaction records, or personnel lists), our task focuses on understanding *visual* information within architectural schemes. Specifically, we aim to detect and interpret the positions and orientations of elements such as windows with respect to a compass. Although LayoutLM could serve as an initial framework for tokenization and information extraction, its architecture is primarily text-centered. Therefore, given the visual nature of our task, we adopt an alternative approach that operates exclusively on visual information.

2.2 Building of Solar radiation exposure evaluation

To the best of the current knowledge, no directly relevant studies were developed to address the specific task of computing solar exposure within architectural schemes, likely due to variations in documentation standards and facility mapping practices across countries. However, several studies have investigated solar exposure estimation using computational and geometric methods.

In particular, the work *Evaluation of Two Solar Radiation Algorithms on 3D City Models for Calculating Photovoltaic Potential* [8] compares two algorithms for computing solar irradiance on three-dimensional urban surfaces. The authors evaluate both a geometric shadow-casting method and a raster-based cumulative radiation model using detailed city-scale 3D models. The results highlight a trade-off between computational efficiency and spatial accuracy: while geometric methods provide higher precision for local surfaces, raster-based approaches allow faster large-scale assessments. Such algorithms are relevant to this project, as they emphasize how 3D modeling and orientation data can significantly improve the estimation of solar exposure on building facades and windows.

2.3 Applications

Among existing open-source solutions capable of interpreting visual information from documents, most are designed for graph, or chart, data extraction rather than complex image analysis. For example, tools like *GraphReader* [9] converts graphical plots into structured datasets, enabling numerical data retrieval from visual charts. However, these systems are primarily focused on extracting quantitative values rather than understanding higher-level spatial or semantic content. Consequently, the present work targets a more specific and underexplored problem—understanding visual representations in architectural schemes—thus requiring

a novel, task-oriented approach.

3 Methodological Approach

In this section we present 3.1) the information setting on which we built our model; 3.2) the data preparation and how it was labeled; 3.3) brief description of the YOLO model and how we used it; 3.4) the information on ResNet18 model how we apply it to our task.

3.1 Information Setting and Steps

The project began with a limited dataset consisting of cadastral scans of residential buildings managed by a local public authority as part of the social housing program. Households living in these buildings were interviewed, and their anwers as well as information from their bills are part of the pilot survey funded by the GRINS' project. Overall we had a 885 scans with different level of quality. Each cadastral scan typically contained:

1. the building address,
2. a schematic layout of the building interior and of the surrounding area,
3. the positions of windows and balconies inside of the buidlding,
4. a compass indicating the north direction (Table 1), and
5. (optionally) the building's position relative to the street.

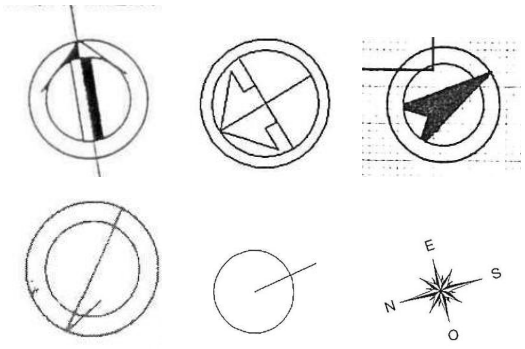


Table 1: Examples of the compass

Given the large dataset size, the need to ensure anonymity, and the goal of minimizing human intervention in future processing, we excluded the option of manual annotation. We proceed by designing an automated pipeline to extract and infer directional information with minimal supervision. The objective was not only to

construct a dataset containing information on building/dwelling orientation for the ongoing analysis on Reggio Emilia residential public housing, but also - more generally - to build a scalable, AI-based system capable of processing large volumes of new scans efficiently.

The proposed model integrates object detection and orientation estimation in a modular structure. Specifically, it operates through four main stages:

1. **Schema and Compass Detection:** A detection model identifies and crops the schematic layout and compass from the full scan.

2. **Window and Balcony Detection:** On the extracted layout, windows and balconies are localized.

3. **Directional Axis Estimation:** The cropped compass image is analyzed to determine the four principal directions (North, East, South, and West, with intermediate diagonals if applicable).

4. **Orientation Mapping:** The extracted directional axes are projected onto the detected layout to determine the orientation of each window and balcony relative to the compass.

This pipeline forms the foundation for generating structured datasets of solar exposure–related attributes from cadastral images. It supports scalability, data privacy, and adaptability across different document types and formats. The modular nature of the approach also allows for future extensions, such as integrating geospatial information when compasses are absent or combining visual features with external mapping data to refine orientation estimation.

3.2 Data preparation and labeling

Before applying any machine learning models, a pre-processing stage is necessary to convert the original dataset from PDF format into standard image formats such as JPEG or PNG. We did such conversion facilitated easier handling and compatibility with deep learning frameworks used in the project.

It is important to note that most scans contain only one schematic layout corresponding to a living area (i.e., not a garage or other auxiliary structures): this characteristic of the scan simplifies the task of identifying the correct layout. In the first step of the pipeline, this layout is typically unique within each scan. Even if there is gonna be a specific living area on the following steps (e.g. windows detection) it will return no results, as it should.
For the detection of schemas, compasses, windows, and balconies (Steps 1 and 2 of

the pipeline), the YOLO object detection model was used. For directional classification from compass images (Step 3), a custom model was implemented using the `ResNet18` architecture from the `PyTorch` library.

The annotation of training data was performed manually using the open-source tool `Label Studio`. Bounding boxes and classification labels were provided for each relevant object (schemas, compasses, windows, and balconies), forming the ground truth required to train and validate the models.

3.3 Detection model - YOLO

For object detection tasks, the YOLO architecture was employed. YOLO is a real-time object detection model known for its speed and accuracy, which predicts bounding boxes and class probabilities directly from full images in a single evaluation. In this project, two separate YOLO models were trained and applied:

- **Schema and Compass Detection:** This model is responsible for detecting and extracting the architectural schemas and compass elements from the input scans. It enables the isolation of relevant content from otherwise complex or noisy documents.
- **Window and Balcony Detection:** Applied to the cropped schema images, this model detects the locations of windows and balconies. These bounding boxes are later used to assign directional labels.

Both models were trained using the Ultralytics implementation of YOLO, which provides a streamlined interface for dataset preparation, training, and inference.

3.4 Classification model - ResNet18

To determine the directional orientation indicated by the compass in each scan, a classification model based on the ResNet18 architecture was used. ResNet18 is a convolutional neural network that utilizes residual connections to enable the training of deeper networks by mitigating the vanishing gradient problem.
In this project, the ResNet18 model processes cropped compass images to extract two key features, namely the position of the compass center and the direction pointing to geographic north.
Using these outputs, the system constructs a set of directional axes (e.g., North, North-East, East, etc.), which are then mapped onto the corresponding building schema to estimate the orientation of detected windows and balconies. Similarly, this approach could be used to estimate the orientation of the dwelling external walls.

4 Results and Discussion

The proposed pipeline achieved high accuracy across all stages. Its evaluation was performed on a held-out validation subset using the Ultralytics implementation of YOLO for object detection and a ResNet-18-based classifier for directional analysis.

Both models were trained on a balanced dataset curated from the annotated scans described in Section 3. Performance metrics were computed following standard detection and classification evaluation protocols.

4.1 YOLO Models

The detection stage consisted of two models: (1) a **Schema & Compass detector** and (2) a **Window & Balcony detector**. Both were fine-tuned from the pretrained YOLOv11 backbone using transfer learning to accelerate convergence and improve accuracy on domain-specific objects. The obtained quantitative results are presented in Table 2.

Precision	**Recall**	**mAP@50**	**mAP@[50:95]**
0.97	0.98	0.99	0.78

Table 2: Results of the YOLO models

Outcomes in Table 2 indicate a strong balance between sensitivity and precision, with minimal false positives. The models correctly localized and classified the majority of schema, windows, and balconies even in scans with variable line thickness and font sizes.

The few detection errors primarily occurred in degraded or low-contrast scans, where the compass was faintly drawn or partially cropped. In some cases, window and balcony symbols were misclassified due to nonstandard annotations or architectural variations (e.g., shaded or labeled windows).

Nevertheless, the models generalized well across unseen layouts, confirming the pipeline's suitability for large-scale automated processing.

4.2 Direction Classification

The orientation analysis component was implemented using a **ResNet-18** Convolutional Neural Network. Its objective was to classify the direction of north and detect the compass center from the cropped compass images produced by the YOLO Schema & Compass detector.

A manually annotated dataset of approximately 900 compass samples was prepared, including various graphical styles (arrows, circles, star-shaped symbols) and rotations. The model was trained for 30 epochs using the Adam optimizer with a learning rate of 5×10^{-3} and categorical cross-entropy loss. The trained classifier achieved:

- **North direction prediction accuracy:** $\sim 98\%$
- **Compass center localization accuracy:** $\sim 96\%$

Qualitative analysis confirmed that the network performed reliably even on compasses with partial occlusions, faded lines, or additional text. Errors typically occurred when the compass was stylized (e.g., decorative labels instead of arrows) or when the north arrow was indistinguishable from other markers.

Here is the example of final work of all models together:

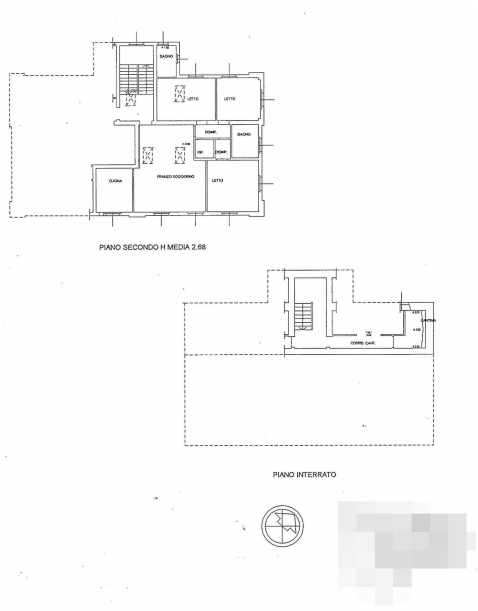
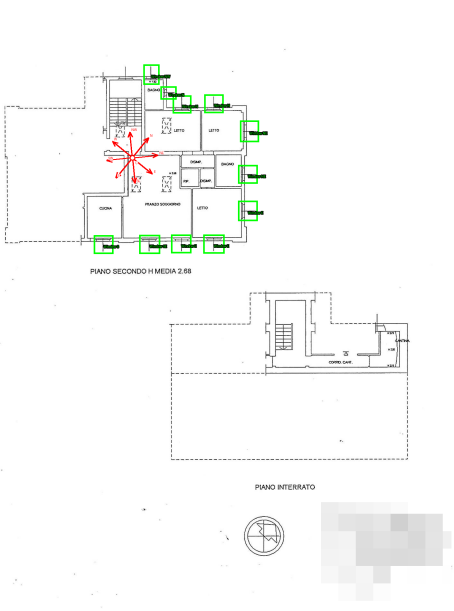
Original Scan	**Processed Output**
	

Table 3: Example of a scan before and after processing by the proposed model

4.3 Output Representation

The final pipeline produces a structured CSV file summarizing extracted information for each processed scan. Each entry contains the detected window and balcony

orientations, counts, potential error codes for the whole scan, and the predominant direction of the living area.

Address	**Win Dir**	**#Win**	**Bal Dir**	**#Bal**	**Err**	**Main Dir**
house_001.jpg	N, NE, E	3	NE	1	0	N
house_002.jpg	S, SW	2	SW, W	2	1	SW
house_003.jpg	E, SE, S	3	-	0	0	SE

Table 4: Example of generated dataset structure.
Notes: *Win Dir* denotes the orientation(s) of windows; *#Win* is the total number of windows per dwelling; *Bal Dir* denotes the orientation(s) of balconies; *#Bal* is the number of balconies; *Err* indicates annotation errors; *Main Dir* represents the dominant facade orientation.

Over the tests, we identify the main possible annotation errors and list them in Table 5.

Error Codes

Code	**Description**
0	No problem
1	No plan and no windows detected
2	Plan detected but compass missing
3	Compass and plan detected but no windows

Table 5: Error code legend

In addition to the detection metrics reported in Table 2, we explicitly monitor an error indicator (*Err*) to capture failure modes that are not fully reflected by aggregate precision–recall statistics. While high mAP values indicate strong average performance of the YOLO models, localized errors such as missed main entities (compass, plan, window, or balcony) or their misclassification may propagate to subsequent stages of the pipeline, affecting main direction classification. For this reason, the error codes reported in Table 5 provide a structured representation of scan-level anomalies and annotation inconsistencies, enabling problematic samples to be flagged. Typical sources of errors, as mentioned earlier, can be reduced by augmenting the training data, extending the annotation schema to cover rare symbol variants, and applying confidence-based thresholds or secondary verification steps for low-confidence detections.

Systematic monitoring of error indicators therefore supports iterative model refinement and improves the robustness and reliability of the extracted orientation statistics used in downstream analyses.

4.4 Future Work

Although the model performs well, I shortly discuss below several improvements that could enhance robustness and applicability.

Model Refinement

Performance could be improved for degraded or inconsistent scans through: (1) preprocessing techniques such as adaptive thresholding or contrast normalization to enhance visual clarity before detection and (2) dataset expansion with more examples of noisy or partially annotated layouts to increase model generalization.

Compass Absence Handling

Some scans lack a compass, currently resulting in flagged errors. Future work could integrate geospatial data to infer orientation in two different ways: (1) using facility addresses with mapping APIs (e.g., OpenStreetMap) to determine building orientation automatically and (2) training a model to estimate directional cues from the geometry of the floor plan itself, leveraging existing compass-annotated data as supervision. These extensions would reduce dependence on explicit compass annotations, extend coverage to more document types, and increase the pipeline's adaptability for real-world large-scale applications.

5 Conclusions

This paper describes a new and automated pipeline we develop to extracting directional and structural information from cadastral building scans. The information gained by the model/pipeline proposed here adds to the comprehensive picture of the socio-economic conditions, energy use, and housing experiences of 250 tenants living in public residential housing in ACER Reggio Emilia which went collected with a dedicated survey. In particular, such information is relevant to interpreting tenants' answers on comfort and to complement existing information on the status of their dwellings.

The proposed model integrates deep learning–based object detection and orientation analysis to identify architectural elements such as windows, balconies, and compasses, and to infer their spatial relationships within the document layout. The work done could be also extended to collect information about walls' orientation of the dwelling considered.

The YOLO-based detectors achieved high precision and recall across different document formats, demonstrating robustness to variations in scale, resolution, and annotation style. The ResNet-18 classifier reliably estimated compass orientation, achieving nearly 98% accuracy and ensuring consistent directional mapping across

diverse visual inputs. Together, these components produced a structured, machine-readable dataset summarizing solar exposure–related features for each processed plan.

Future work will focus on extending the system's adaptability to cases where the compass is absent, integrating geospatial data for orientation inference, and enhancing model performance on degraded or noisy scans. These developments will further expand the applicability of the method to large-scale, real-world building datasets and support downstream analyses such as photovoltaic potential estimation and daylight modeling.

Overall, we confirm that visual information in architectural documents can be automatically interpreted with minimal human supervision, enabling scalable data collection for energy efficiency and solar analysis applications. The proposed approach also preserves data privacy, as it operates entirely on visual inputs without relying on textual content.

Moreover, the solar orientation information extracted and systematically organized by the proposed model is highly relevant for improving energy efficiency in the residential sector. Accurate knowledge of window and balcony orientation enables a more precise assessment of solar gains, daylight availability, and thermal performance, which are key factors in both building design and retrofit planning. By transforming unstructured cadastral scans into machine-readable datasets, the proposed approach facilitates large-scale analyses that were previously impractical due to data and resource constraints.

In this perspective, future research building on the present work may provide a robust empirical basis for the design of evidence-based energy policies, regulatory frameworks, and targeted incentive schemes aimed at reducing energy consumption and mitigating energy poverty. Moreover, such information can support the development of new methodologies for optimizing building renovation strategies, promoting the adoption of passive solar design principles, and improving the allocation of public resources. Ultimately, the integration of automated visual data extraction with energy and socio-economic analyses has the potential to contribute to more sustainable environmental outcomes and to enhance the economic efficiency and resilience of housing systems at both national and regional levels.

Acknowledgments

The author acknowledges the financial support from the European Union - NextGenerationEU, in the framework of GRINS - Growing Resilient, INclusive and Sustainable project (GRINS PE00000018 - CUP C93C22005270001) within Spoke 6 *Low Carbon Polices* activities. The views and opinions expressed are solely those of the author and do not necessarily reflect those of the European Union, nor can the

European Union be held responsible for them. Special thanks to Marta Castellini, Benedetta Grassi and Paola Valbonesi for insightful discussions on the development of the model/pipeline and its application to the investigated case of ACER Reggio Emilia Residential Public Housing.

References

[1] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," 2016. [Online]. Available: https://arxiv.org/abs/1506.02640

[2] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," 2015. [Online]. Available: https://arxiv.org/abs/1512.03385

[3] Y. Xu, M. Li, L. Cui, S. Huang, F. Wei, and M. Zhou, "Layoutlm: Pre-training of text and layout for document image understanding," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery amp; Data Mining*, ser. KDD '20. ACM, Aug. 2020, p. 1192–1200. [Online]. Available: http://dx.doi.org/10.1145/3394486.3403172

[4] S. Mishra, A. Gauquier, and P. Senellart, "Modular multimodal machine learning for extraction of theorems and proofs in long scientific documents (extended version)," 2024. [Online]. Available: https://arxiv.org/abs/2307.09047

[5] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," 2019. [Online]. Available: https://arxiv.org/abs/1810.04805

[6] G. Jaume, H. K. Ekenel, and J.-P. Thiran, "Funsd: A dataset for form understanding in noisy scanned documents," 2019. [Online]. Available: https://arxiv.org/abs/1905.13538

[7] A. W. Harley, A. Ufkes, and K. G. Derpanis, "Evaluation of deep convolutional nets for document image classification," in *2015 13th International Conference on Document Analysis and Recognition (ICDAR)*. IEEE, 2015, pp. 991–995.

[8] S. M. Murshed, A. Simons, A. Lindsay, S. Picard, and C. D. Pin, "Evaluation of two solar radiation algorithms on 3d city models for calculating photovoltaic potential," in *Proceedings of the 4th International Conference on Geographical Information Systems Theory, Applications and Management (GISTAM), Volume 1*. SciTePress, 2018, pp. 296–303.

[9] "Graphreader – online tool for reading graph image values and save as csv/json," https://www.graphreader.com/, 2025, accessed: 2025-10-31.