



**Politecnico
di Torino**

Polytechnic University of Turin
Master's Degree in Green Buildings Engineering
A.Y. 2025/2026

**BIM and Construction Project Management for Photovoltaic Systems
in Industrial Buildings**

Case Study: Stellantis factory located in San Benigno Canavese, Italy.

Supervisor:
Prof. Anna Osello

Co-Supervisor:
Prof. Matteo Del Giudice

Candidate:
Amir Alebrahimrovies

Acknowledgments

First and foremost, I would like to express my deepest gratitude to my esteemed supervisor, **Prof. Anna Osello**, for her invaluable guidance, immense support, and continuous encouragement throughout the development of this thesis. Her profound expertise, insightful critiques, and inspiring vision have not only shaped this research but have also profoundly enriched my academic and professional growth.

I am equally grateful to **Prof. Matteo Del Giudice** for his exceptional support and valuable insights. His constructive feedback, dedicated time, and expertise were instrumental in overcoming the challenges encountered during this work and significantly contributed to the successful completion of this project.

Behind every academic achievement is a foundation of love and unwavering support. I owe a debt of gratitude to my family, whose endless encouragement, patience, and belief in my abilities have been my greatest source of strength during this long journey.

A very special and heartfelt thanks goes to my little brother, **Saeed**. More than just a brother, he has been my ultimate support system—acting as my primary "sponsor" and adding an essential element of fun, humor, and joy to my life when the academic pressure seemed overwhelming. This milestone would not have been the same without him.

To everyone who contributed directly or indirectly to this milestone, thank you for being a part of this journey.

Note on AI Assistance: In the preparation of this thesis, I have utilized artificial intelligence tools to assist with language refinement and editing. I confirm that the originality, accuracy, research methodology, and final results presented in this work are entirely my own.

Abstract

Industrial rooftops offer one of the most underutilised opportunities in commercial decarbonisation, yet photovoltaic retrofit conversion rates remain low because structural assessment, solar simulation, layout design, cost estimation, and construction planning each occupy isolated software environments. Synchronising them demands constant manual re-entry, an overhead referred to throughout this thesis as the “iteration penalty.” This study develops and validates an integrated BIM-centred framework that eliminates this penalty by consolidating geometry extraction, solar ray-tracing, multi-criteria financial optimisation, construction project management scheduling, and 7D semantic embedding within a single uninterrupted Autodesk Revit session, with no file export required at any stage. Spatial geometry is extracted through Dynamo and passed to an embedded CPython engine that executes full-year reverse ray-tracing over a Turin-Caselle EnergyPlus Weather file, filtered from 8,760 to 4,388 active daylight hours via a dayMask pre-processing step. The framework is validated on the Stellantis automotive manufacturing facility at San Benigno Canavese, Piedmont, Italy: 26,372 m² of gross roof area, 8,460 virtual sensor points, a peak annual irradiance of 1,891 kWh/m²/year on south-facing sawtooth surfaces, and a site-wide weighted average electrical yield of 169.03 kWh/m²/year. Cross-validation against Ladybug Tools confirms agreement within 1.5 %. A multi-criteria decision analysis evaluates 135 configurations ($3 \times 3 \times 3 \times 5$: module wattage, layout topology, inverter type, battery storage) and selects an 8,390-module, 3.78 MWp string-inverter system as Rank 1, generating 3.12 GWh per year and offsetting 20.8 % of the facility’s 15 GWh annual demand. Total capital expenditure, priced against the Prezziario Regionale Piemonte, is €1,380,325, with a payback period of 3.17 years, an annual return on investment of 31.52 %, and a projected construction duration of 18 months requiring approximately 10,983 labour hours. All of these figures are derived and updated automatically within the BIM environment as design decisions evolve. Over a 25-year horizon the system avoids approximately 33,500 tCO₂e. The principal contribution is methodological: demonstrating that ray-tracing, financial optimisation, construction project management, and IFC parameter write-back can be executed natively within a single BIM session, producing a semantically rich digital twin that supports both capital investment decisions and long-term operational management without disciplinary discontinuity.

Keywords: Building Information Modeling; photovoltaic retrofit; industrial rooftop solar; construction project management; Dynamo visual programming; CPython ray-tracing; EnergyPlus Weather; multi-criteria decision analysis; 7D digital twin; IFC parameterisation; Revit API

Contents

Acknowledgments.....	2
Abstract.....	3
Figures List:	8
Tables List:.....	11
CHAPTER 1	12
Introduction.....	12
1.1 Background.....	13
1.1.1 Energy Retrofitting in the Industrial Sector.....	15
1.1.2 Spatial and Structural Challenges of Industrial Roofs	17
1.1.3 Technical Constraints: Shading and Electrical Yield	20
1.2 Building Information Modeling (BIM): Fundamentals and Evolution.....	22
1.2.1 BIM Maturity and Multidimensionality.....	23
1.2.2 Level of Development (LOD) and Level of Information (LOI)	25
1.3 The Paradigm Shift: BIM Beyond Geometry	26
1.4 Computational Interoperability and Workflow Efficiency	31
1.4.1 The Bottleneck of File-Based Exchange and the Iteration Penalty	32
1.5 Problem Statement and Industry Bottlenecks	35
1.6 Research Objectives and Methodological Framework	39
1.7 Scope and Limitations of the Study	40
1.8 Research Questions	41
1.9 Chapter Summary	42
CHAPTER 2	44
Literature Review.....	44
2.1 Introduction to the Literature	44
2.1.1 Purpose and Scope of the Review.....	45
2.1.2 Review Methodology and Search Strategy	46
2.2 Algorithmic Spatial Coordination and Visual Programming.....	47
2.2.1 Case Study: BIPV Automated Layout Optimisation in Industrial Environments ..	48
2.3 4D BIM: Construction Sequencing and Site Logistics	49
2.3.1 Case Study: Logistical Phasing in Congested, Zero-Downtime Environments	50
2.4 5D BIM: Cost Integration and Live Financial Estimation.....	51
2.5 7D BIM: Lifecycle Data and Digital Twin Integration	52
2.6 The Integrated Research Gap: Synthesis	53
2.7 Literature Matrices and Summary	54
2.8 Chapter Summary	59
	5

CHAPTER 3	61
Methodology and Computational Framework	61
3.1 Introduction and System Architecture	61
3.1.1 The Tool Ecosystem: A Tripartite Architecture	61
3.1.2 The Continuous Feedback Loop	64
3.2 Case Study Context: The Stellantis San Benigno Canavese Facility	66
3.2.1 Geographical and Environmental Profile.....	67
3.2.2 Architectural Complexity and Existing Structural Topology	68
3.2.3 Digitisation of the BIM Baseline	70
3.2.4 Energetic Baseline and Managerial Directives	71
3.3 Algorithmic Spatial Discretisation: The Geometric Engine	72
3.3.1 Native Extraction and Topological Vector Filtering	72
3.3.2 Sensor Grid Generation via UV Domain Division	73
3.3.3 Area Estimation and Packing Factor Calibration.....	75
3.4 Native Environmental Simulation and Ray-Tracing: The Physics Engine.....	76
3.4.1 Meteorological Integration and Computational Optimisation	76
3.4.2 Reverse Ray-Tracing Execution and Raw Irradiance Calculation	78
3.4.3 Dynamic Heat Mapping and Visual Feedback	79
3.5 5D Financial Optimisation and Scenario Modelling: The Cost Engine	80
3.5.1 Parametric System Variables and Yield Normalisation.....	81
3.5.2 System Sizing: Resolving Spatial versus Energetic Constraints	82
3.5.3 Multi-Criteria Scenario Evaluation and CAPEX Optimisation	83
3.6 Automated Layout Generation and 7D Semantic Embedding: The Output Engine	86
3.6.1 Methodological Comparison of Layout Generation Techniques.....	86
3.6.2 Algorithmic Instantiation of the Optimised Scenario	88
3.6.3 7D Semantic Embedding and IFC Parameterisation	88
3.7 Methodological Validation via Cross-Platform Analysis	89
3.7.1 Geometric Export and Parameter Alignment.....	89
3.7.2 Comparative Incident Radiation Analysis	90
3.7.3 Conclusion of the Computational Methodology	92
CHAPTER 4	93
Results and Discussion	93
4.1 Introduction to the Findings.....	93
4.1.1 Methodology Findings	93
4.2 Technology and Technical Findings	95
4.3 Interoperability Findings.....	99

4.4 Energy and Sustainability Findings	101
4.5 Research Question Synthesis	105
4.6 Consolidated Results.....	107
4.7 Chapter Summary	110
CHAPTER 5	112
Conclusion and Future Directions	112
5.1 General Conclusion.....	112
5.2 Limitations of the Current Study	113
5.3 Future Possibilities and Research Directions.....	114
5.3.1 Integration of Hybrid Solar Systems (PV + Thermal).....	114
5.3.2 Heat Pump Integration and COP Optimisation.....	116
5.3.3 Advanced PV Technologies and BIPV Integration	116
5.3.4 Integration of Biomass for Baseload Supplementation	117
References.....	118

Figures List:

Figure 1 Conceptual workflow comparison: API-driven BIM methodology vs. traditional project delivery. Source: Author's elaboration, adapted from BIM-based integrated project delivery literature.	13
Figure 2 Greenhouse gas emissions from industrial processes for the top ten emitting countries (1990–2020). Source: [1].	14
Figure 3 Historical and projected global renewable electricity capacity additions by technology and segment (2016–2028). Source: IRENA [2].	15
Figure 4 EU ETS carbon allowance prices alongside the Levelised Cost of Electricity (LCOE) of utility-scale solar PV. Source: Author's elaboration based on data from [1] and [2].	17
Figure 5 Spatial and structural limitations of industrial roofs, detailing restricted zones for photovoltaic installations. Source: Author's elaboration.	18
Figure 6 Severity and impact of partial shading degradation on relative power output within a PV series string. Source: Author's elaboration.	21
Figure 7 Effect of partial shading on a series string and its corresponding electrical mismatch principles. Source: Author's elaboration based on Woyte et al. [5].	22
Figure 8 Historical Building Information Modeling (BIM) development timeline. Source: Adapted from Eastman et al. [9].	23
Figure 9 The Bew-Richards BIM Maturity Model. Source: Adapted from Bew & Richards [11].	24
Figure 10 Dimensions of Building Information Modeling. Source: Adapted from Succar [12].	25
Figure 11 Geometric progression from conceptual representation (LOD 100) to fabrication (LOD 400). Source: Author's elaboration based on AIA specifications [13].	27
Figure 12 Progression of non-graphical data within the proposed BIM framework corresponding to Level of Information (LOI) tiers. Source: Author's elaboration.	27
Figure 13 System architecture of the proposed API-driven multidimensional workflow, centralising 4D, 5D, and 7D data distribution. Source: Author's elaboration.	28
Figure 14 Visual programming workflow in Dynamo for extracting roof boundaries, generating spatial grids, and filtering panel coordinates against MEP obstacles. Source: Author's elaboration.	29
Figure 15 Embedded Python scripting within the Dynamo visual programming environment mapping the spatial discretisation matrix. Source: Author's elaboration.	30
Figure 16 System architecture of the continuous loop methodology using Dynamo data routing and the embedded Python physics engine. Source: Author's elaboration.	30
Figure 17 API-driven BIM workflow integrating Revit, Dynamo, and Python for automated design coordination, simulation, and project management updates. Source: Author's elaboration.	32
Figure 18 Limitations of traditional file-based interoperability workflows (IFC and gbXML) resulting in manual updates and workflow inefficiency. Source: Author's elaboration based on Dong et al. [23] and Venugopal et al. [24].	33
Figure 19 Proposed API-driven BIM workflow integrating Revit, Dynamo, Python, and climate data within a unified computational environment. Source: Author's elaboration.	34
Figure 20 Sequence diagram illustrating the latency and rework cycles of traditional multidisciplinary communication. Source: Author's elaboration.	36
Figure 21 Multidimensional workflow separation in conventional BIM workflows illustrating how isolated software tools result in the Iteration Penalty. Source: Author's elaboration.	37

Figure 22 Radar chart mapping the performance of existing computational workflows against the proposed native API-based methodology. Source: Author's elaboration.	38
Figure 23 Chapter 1 conceptual overview. Source: Author's elaboration.	43
Figure 24 Comparison of the legacy fragmented workflow and the proposed unified API-driven loop, illustrating how continuous data integration eliminates the Iteration Penalty in industrial PV retrofit project management. Source: Author's elaboration.	45
Figure 25 Data exchange architecture showing bidirectional communication between Autodesk Revit, Dynamo, and an embedded Python engine integrated with external climate and regional cost data. Source: Author's elaboration.	64
Figure 26 Aerial context of the Stellantis manufacturing facility in San Benigno Canavese, Piedmont.	67
Figure 27 Master plan and architectural site layout of the industrial building complex used as the case study for the rooftop PV retrofit project.	68
Figure 28 Aerial satellite view of the industrial plant site, highlighting the 25,000 m ² rooftop footprint evaluated for the photovoltaic retrofit installation.	69
Figure 29 Breakdown of the federated BIM model into its constituent electrical, mechanical, and architectural sub-models for coordinated engineering analysis. Source: BIM-to-BEM Framework for Energy Retrofit in Industrial Buildings: From Simulation Scenarios to Decision Support Dashboards.	69
Figure 30 Extracted roof category from Revit LoD300 Model	70
Figure 31 Dynamo visual programming workflow for automated host surface filtering, demonstrating the extraction of structural roof faces and the subsequent flattening of surface data lists for solar grid mapping.	73
Figure 32 Dynamo algorithm executing vector dot product calculations to evaluate surface normals against the Z-axis, filtering out unsuitable faces into a refined group of valid host surfaces (facesUp).	73
Figure 33 Dynamo visual programming workflow dividing the facesUp surfaces into UV sensor grids.	74
Figure 34 Visualisation of the generated sensor grid layout (PtsFlat) distributed uniformly across the selected industrial roof plane for high-resolution solar irradiance sampling.	74
Figure 35 Plan view of the full industrial rooftop layout, showing the comprehensive grid distribution of the 8,460 generated virtual sensor points across the selected structural planes.	75
Figure 36 Dynamo workflow for rooftop area estimation and packing factor application.	76
Figure 37 Dynamo and Python script configuration parsing the external EnergyPlus Weather file and applying the dayMask pre-processing filter.	77
Figure 38 Computational pipeline showing the reduction from 8,760 to 4,388 active simulation hours after dayMask application.	77
Figure 39 Dynamo visual programming logic calculating the total roof solar potential from aggregated sensor-point irradiance values.	79
Figure 40 Dynamo visual programming pipeline for automated heat map generation, showing the normalisation routine and colour gradient assignment.	80
Figure 41 Dynamo visual programming interface displaying the 'Project & Panel Inputs' code block.	81
Figure 42 Dynamo visual programming logic executing the dynamic yield normalisation process.	82
Figure 43 Dynamo visual programming logic computing the area-weighted average electrical yield.	82

Figure 44 Dynamo visual programming logic determining the theoretical system capacity based on the facility's baseline energetic demand.....	83
Figure 45 Dynamo visual programming logic executing the area limitation check and conflict resolution.....	83
Figure 46 Dynamo visual programming logic evaluating final metrics for the maximised system density.....	83
Figure 47 Dynamo visual programming interface displaying the output of the Python optimisation script.....	86
Figure 48 Multi-Criteria Decision Analysis (MCDA) matrix evaluating the top five system configurations.	86
Figure 49 Cross-platform validation setup in Rhinoceros 3D and Ladybug Tools showing the imported industrial roof geometry, solar analysis grid, and sky-dome environment based on the Turin-Caselle EPW dataset.	90
Figure 50 Grasshopper–Ladybug incident radiation analysis workflow used for validating the custom Dynamo–Python solar simulation framework.....	91
Figure 51 (a) Annual ROI comparison across the six representative configurations from the 135-scenario MCDA matrix. (b) Payback period comparison. (c) CAPEX breakdown by system component for the Rank 1 configuration.	95
Figure 52 (a) Frequency distribution of annual irradiance values across all 8,460 virtual sensor points. (b) Roof area breakdown from gross geometric to net constructible area after packing factor and obstacle exclusion. Source: Author’s elaboration.	97
Figure 53 a) 25-year cumulative net cash flow for Rank 1 (b) Monthly generation profile for the 3.78 MWp Wide configuration at Turin latitude, derived from TMY3 EPW dataset	103
Figure 54 a) 25-year cumulative carbon avoidance trajectory (b) Annual energy waterfall diagram.....	105
Figure 55 Radar chart comparing the proposed API-native integrated framework with representative conventional workflows across the four research dimensions.	107
Figure 56 Integrated BIM-centred workflow consolidating geometry extraction, solar simulation, financial optimisation, and construction project management within a single Revit session. Source: Author’s elaboration.....	112
Figure 57 Hybrid system integration with facility demand	115
Figure 58 Conversion efficiency comparison of PV and Solar Thermal.....	115
Figure 59 Conceptual integration of rooftop PV generation with an industrial heat pump, showing the conversion of surplus solar electricity into stored thermal energy for process heat and space heating.	116
Figure 60 Comparison of monocrystalline PV and thin-film BIPV, highlighting expanded envelope coverage and greater decarbonisation potential.	117

Tables List:

Table 1 Categorisation of structural loading constraints in industrial photovoltaic retrofits. Source: Author's elaboration based on EN 1991 standards [4].	19
Table 2 Shadow length variation for a 2-metre rooftop obstacle at solar noon (45°N). Source: Author's elaboration.	20
Table 3 Master paper feature audit: all reviewed studies mapped against integration dimensions.	55
Table 4 Capability and gap summary by thematic cluster.	57
Table 5 Integration maturity assessment of key studies.	58
Table 6 Binary cross-dimensional audit. ✓ = present; ✗ = absent; ● = partial. Source: Author's elaboration.	58
Table 7 Summary of tools, their roles and design rationale within the tripartite computational architecture. Source: Author's elaboration.	62
Table 8 Complete workflow stage summary showing tools, outputs, methods, and significance for each of the eleven computational stages. Source: Author's elaboration.	65
Table 9 MCDA results for the 135 configurations evaluated. ★ = selected Rank 1 configuration. Source: Author's elaboration based on Prezziario Regionale Piemonte [57]...	84
Table 10 Comparative assessment of PV layout generation methods. ✓ = capability present; ✗ = absent. Source: Author's elaboration.	87
Table 11 Parameter-by-parameter cross-platform validation comparison. Source: Author's elaboration.	90
Table 12 Key Performance Indicators: Rank 1 (Optimal) Configuration.	94
Table 13 Spatial Optimisation: Theoretical Maximum vs. Optimised Rank 1 Configuration	96
Table 14 Multi-criteria decision analysis results	98
Table 15 String vs. Central Inverter Architecture: Technical Comparison	99
Table 16 7D Semantic Data Layer: Parameters, Storage Mechanisms, and Operational Use Cases	100
Table 17 Financial and Operational Performance: Rank 1 Configuration and Full Scenario Range	102
Table 18 Energetic Offset and Environmental Impact: Complete Outcome Profile	104
Table 19 Research Question Response Matrix: Findings, Evidence, and Practical Significance	109
Table 20 Consolidated Results and Design Implications Matrix.	110
Table 21 Limitation Impact Matrix.	114

CHAPTER 1

Introduction

Industrial rooftop solar seems, on the surface, like a solved problem. The roof is already there. Photovoltaic technology is commercially mature and well understood. And after a decade of dramatic cost reductions, the economics are compelling for most industrial applications. So why do so many of these projects still take longer than planned, cost more than budgeted, and deliver less energy than originally estimated?

The hardware is rarely to blame. Modern panels and inverters are reliable, predictable, and increasingly affordable. The difficulties almost always trace back to something less visible and more frustrating: the fragmented, tool-hopping process through which industrial PV retrofits are typically planned and coordinated. Structural checks happen in one software environment, energy simulations in another, cost estimates in spreadsheets, and schedules in yet another platform entirely. None of these systems talk to each other automatically. So when the design changes, someone has to manually chase the update across every tool. That coordination burden is the real source of delay and budget creep.

The interdependence of design decisions makes this especially acute. Moving a row of panels by a few metres to satisfy a structural constraint sounds trivial. In practice, that single adjustment can alter shading across neighbouring arrays, change cable lengths, shift material quantities, redistribute structural loads, and rewrite the project's financial projections. Every decision propagates. In a workflow where tools are disconnected, none of those knock-on effects update themselves. A human being has to find them, check them, and manually reconcile the datasets, often multiple times as the design evolves.

This thesis proposes a different approach. Rather than treating geometric modelling, energy simulation, cost analysis, and scheduling as separate activities connected by file exports, the methodology developed here maintains all of these within a single computational framework built around Autodesk Revit. The Revit API, combined with Dynamo visual programming and embedded Python scripting, allows design changes to propagate automatically through performance calculations, cost models, and scheduling outputs. No manual file transfer. No periodic reconciliation. Information stays consistent because it was never disconnected in the first place.

The goal is not to produce a perfect simulation of every physical phenomenon on an industrial rooftop. It is, more practically, to eliminate the administrative overhead that currently prevents engineers from exploring more than a handful of design alternatives before a deadline forces a decision. Figure 1 places this objective in context, comparing the conventional fragmented workflow with the integrated API-driven approach developed in this research.




Figure 1 Conceptual workflow comparison: API-driven BIM methodology vs. traditional project delivery. Source: Author's elaboration, adapted from BIM-based integrated project delivery literature.

1.1 Background

Industrial interest in rooftop photovoltaics is not new. Manufacturers have been exploring on-site solar generation for well over a decade, drawn by the prospect of lower operating costs and greater energy independence. What has genuinely changed in recent years is the speed of adoption and the two forces driving it have almost nothing to do with each other.

The first is regulatory and reputational pressure to decarbonise. Manufacturing remains one of the most energy-intensive sectors in the global economy, and companies now face persistent scrutiny from regulators, investors, and customers who want measurable progress on emissions

reduction. Declaring sustainability ambitions is no longer enough; demonstrating them with actual infrastructure is increasingly expected.

The second force is purely economic. The cost of photovoltaic technology has fallen dramatically by roughly 90 per cent over the past decade, and that shift has transformed rooftop solar from a corporate responsibility gesture into a straightforward infrastructure investment with a calculable return. For most industrial facilities today, the question is no longer whether PV makes technical sense. It is whether there is a good reason not to use available roof space for electricity generation.

Figure 2 puts the emissions context in perspective. Industrial process emissions for the ten highest-emitting countries between 1990 and 2020 collectively exceeded 2,000 MtCO₂e [1]. The trajectory, however, looks very different depending on which country you examine. China's industrial emissions rose sharply from the early 2000s and eventually dwarfed those of every other economy. The United States and the European Union, by contrast, remained comparatively stable. This reflects structural shifts in industrial activity, efficiency improvements, and the gradual tightening of carbon regulation [1].

The practical implication for manufacturers is this: the financial case for on-site renewable generation is stronger now than it has been at any previous point, and the regulatory environment is only moving in one direction. Rooftop photovoltaics happen to address both pressures simultaneously, which explains why they have moved from the periphery of industrial energy strategy to something closer to the mainstream.




Figure 2 Greenhouse gas emissions from industrial processes for the top ten emitting countries (1990–2020). Source: [1].

Figure 3 shows where renewable energy deployment is headed. Solar PV and wind are expected to account for between 90 and 95 per cent of all newly installed renewable capacity through 2028 [2]. Under the most accelerated scenarios, annual additions could exceed 900 GW, with solar contributing roughly 700 GW in a single year. These are not projections from the edges of the forecast distribution. They are the central scenarios from major energy agencies, and investment patterns are already consistent with them [2].

For industrial facility owners, the relevance is straightforward. Manufacturing sites typically have large, flat, underused roofs – structural assets that happen to be well suited to photovoltaic installation. Generating electricity on-site reduces exposure to grid price volatility, contributes to carbon reduction targets, and, over a 25-year system lifetime, tends to look very attractive against the alternative of simply buying power from the grid.




Figure 3 Historical and projected global renewable electricity capacity additions by technology and segment (2016–2028).
Source: IRENA [2].

1.1.1 Energy Retrofitting in the Industrial Sector

Industrial facilities are not easy buildings to manage from an energy perspective. Unlike offices or residential buildings, they cannot simply shift their demand to cheaper time periods. Production equipment runs when production runs. Compressed-air systems, climate control, lighting, and ancillary infrastructure operate continuously, and the electricity demand profile

that results is large, relatively stable, and largely inflexible – it doesn't give facility managers much room to manoeuvre when prices spike.

For energy-intensive sectors like automotive manufacturing, metal processing, and chemicals, this is not a minor inconvenience. Electricity costs represent a significant fraction of operating expenditure. Sustained increases in grid prices translate directly into reduced competitiveness, and that pressure has become harder to ignore as European energy markets have grown more volatile.

Piedmont, where the case-study facility examined in this thesis is located, illustrates this dynamic well. The region has historically been one of Italy's principal industrial centres, with a dense concentration of manufacturing across multiple sectors. Affordable energy was, for many years, something that could be taken for granted. That is no longer the case. Rising price volatility has pushed manufacturers throughout the region to reassess their relationship with the grid and to look more seriously at on-site generation as a structural response rather than a stopgap measure.

Across Europe, two developments have reinforced each other in ways that make rooftop solar particularly attractive right now. First, volatile gas prices have increased the premium on self-generated electricity, especially for facilities running substantial daytime loads. Second, the EU Emissions Trading System has steadily raised the cost of carbon-intensive energy, making the relative economics of low-carbon alternatives more favourable with each passing year. By 2024, these forces were not acting independently; they were compounding. The financial, regulatory, and reputational arguments for on-site renewable generation were all pointing in the same direction simultaneously [1].

Figure 4 makes this convergence tangible. Between 2014 and 2024, EU ETS carbon allowance prices climbed steadily as emission caps tightened. Over exactly the same period, the levelised cost of solar electricity fell by approximately 90 per cent [1,2]. The lines cross in a way that renders the old calculation (grid power cheap, solar expensive) simply obsolete.




Figure 4 EU ETS carbon allowance prices alongside the Levelised Cost of Electricity (LCOE) of utility-scale solar PV.
Source: Author's elaboration based on data from [1] and [2].

None of this means that every industrial PV retrofit is straightforward to implement. Complete energy independence is rarely achievable for a facility running continuous production at scale. But reducing purchased electricity by a meaningful margin 20, 30, or 40 per cent depending on the roof area and the facility's load profile is often within reach, and when the system is designed properly, it typically delivers returns over a 25-year lifetime that are difficult to match through other capital investments.

The hard part is not the technology. The hard part is the design and coordination process. Available roof area, structural load-bearing capacity, rooftop equipment layout, orientation, and operational constraints all interact in ways that require multiple engineering disciplines working in parallel. And those disciplines, in current practice, rarely use the same tools or share information in real time. That is the problem this research is designed to address.

1.1.2 Spatial and Structural Challenges of Industrial Roofs

One of the most persistent misconceptions in early-stage PV planning is that the total roof area of an industrial building gives a reasonable first estimate of what is actually available for solar panels. It does not. The gap between gross roof area and genuinely usable installation space is routinely large, sometimes exceeding 50 per cent, a gap consistently reported in the wider solar-potential-modelling literature [8], and underestimating it at the feasibility stage has a way of creating painful surprises later.

Industrial rooftops are working surfaces. They carry HVAC equipment, ventilation arrays, exhaust stacks, skylights, cable trays, drainage components, maintenance walkways, and often

process-related pipework that was installed decades before anyone thought about integrating solar panels [3]. None of this equipment was positioned with future PV installation in mind. It occupies physical space, creates access restrictions, and critically casts shadows. The shadows matter enormously because they change throughout the day and across the seasons, meaning that areas which look perfectly clear on a two-dimensional plan can turn out to be heavily shaded in practice.

A concrete example illustrates the scale of the problem. A roof offering 3,000 m² of gross area might ultimately yield fewer than 1,500 m² of viable installation space once equipment clearances, maintenance corridors, fire-safety setbacks, and realistic shading exclusions are applied. That is not an edge case, it is a common outcome on older industrial facilities. Without a detailed three-dimensional assessment, it is very easy to start a project with optimistic capacity projections and arrive at detailed design with a significantly smaller system than the business case assumed.

The geometry of industrial buildings compounds the difficulty. Many facilities are not simple, single-level boxes with uniform flat roofs. They have multiple roof levels, varying pitches, sawtooth profiles designed to capture daylight, and structural bays with different spans and capacities [3]. In practice, the ‘roof’ is often several different zones with different orientations, different shading environments, and different structural characteristics. Treating them as a single homogeneous surface, as simplified area-based feasibility methods do, introduces errors that only become apparent once engineering analysis begins. Figure 5 illustrates some of the constraint types commonly encountered.




Figure 5 Spatial and structural limitations of industrial roofs, detailing restricted zones for photovoltaic installations.
Source: Author's elaboration.

Then there is the structural question, which tends to receive less attention in early planning than it deserves. Many industrial buildings in active use today were designed and constructed decades before rooftop PV became a realistic consideration. Their structural frameworks were sized to resist snow and wind loads as specified in Eurocode EN 1991 at the time of construction [4]. The residual capacity available for additional permanent loads panels, racking, ballast, inverters, cable trays varies considerably from building to building and from zone to zone within the same building. Assuming that residual capacity exists without actually checking is a risk that has caught out more than a few projects.

Table 1 Categorisation of structural loading constraints in industrial photovoltaic retrofits. Source: Author's elaboration based on EN 1991 standards [4].

Primary Load Category	Specific Load Element	Governing Standard	Engineering Impact on Industrial PV Retrofits
Environmental (Variable)	Snow Loads	Eurocode EN 1991-1-3 [4]	Establishes the baseline structural capacity. PV panels alter the roof's thermal profile and physical geometry, which can trap snow and create localised drifting loads that exceed the original design limits.
Environmental (Variable)	Wind Actions	Eurocode EN 1991-1-4 [4]	Generates aerodynamic uplift and lateral sliding forces against the tilted PV arrays. Directly dictates the required mass of the ballast system or the tensile strength of mechanical fasteners.
System Imposed (Permanent)	Modules & Racking	(Dead Load)	Introduces a continuous, distributed dead load (typically measured in kg/m ²) across the roof membrane, reducing the residual structural capacity.
System Imposed (Permanent)	Ballast Systems	(Dead Load)	To avoid penetrating the roof membrane, flat industrial roofs rely on heavy concrete ballasts. These introduce severe, localised point-loads that frequently become the limiting structural factor for older buildings.
System Imposed (Permanent)	Electrical Infrastructure	(Dead Load)	Introduces localised concentrated loads, primarily from heavy commercial string inverters, main distribution boards, and extensive cable tray routing.

What makes this structurally challenging is not any individual load category in isolation, it is the combination. Distributed panel weight, concentrated ballast loads, localised inverter masses, and the interaction of all of these with pre-existing snow and wind demands must be assessed simultaneously. Change one element and the others shift. Relocate a panel row to reduce loading in one zone, and you may create a new problem in an adjacent one. This interdependence is precisely why sequential, discipline-by-discipline workflows struggle on

projects of this kind and why the three-dimensional, parametric approach developed in this research offers a more effective alternative.

1.1.3 Technical Constraints: Shading and Electrical Yield

Shading looks simple. An obstacle blocks sunlight; the panel behind it produces less electricity. But the reality is considerably more complicated, and the gap between intuition and actual system behaviour has caused more than a few PV installations to underperform.

The first complexity is geometric. The sun does not stay in one place. At 45°N latitude (the approximate location of the facility examined in this thesis), the solar elevation angle at noon ranges from 21.35° in December to 68.25° in June. That range means a fixed rooftop obstacle of two meters' height will cast a shadow of about 0.80 m on a summer afternoon and 5.12 m on a winter midday. The same piece of HVAC equipment that barely disturbs production in July can shade several entire panel rows in December.

Table 2 Shadow length variation for a 2-metre rooftop obstacle at solar noon (45°N). Source: Author's elaboration.

Date (Solar Noon)	Solar Elevation Angle (α)	Shadow Length Multiplier	Shadow Length (L)	Impact on PV Area
Summer Solstice (June 21)	68.25°	$0.40 \times h$	0.80 m	Minimal impact
Autumn Equinox (Sept 21)	44.80°	$1.01 \times h$	2.02 m	Moderate impact
Winter Solstice (Dec 21)	21.35°	$2.56 \times h$	5.12 m	Severe impact

The design implications are awkward. If you dimension panel spacing to avoid inter-row shading in summer, you will almost certainly underestimate how bad it gets in winter. If you design for worst-case December geometry, you will leave a lot of usable summer roof space empty. Neither extreme is right. Finding the appropriate balance requires actually running the numbers for the full year not estimating from a single seasonal snapshot.

The second complexity is electrical, and it is genuinely non-intuitive. In a conventional string-wired PV system, all modules in a string carry the same current. The weakest module dictates the performance of the entire string. This means that partial shading on even a single panel say, the shadow of a ventilation stack crossing one corner of one module can drag down the output of every other panel connected in that string [5]. The shaded area might represent 2 or 3 per cent of the array surface. The power loss can be an order of magnitude larger.

Modern modules include bypass diodes that partially mitigate this current routes around shaded cell groups, limiting hotspot risk and protecting the module [6,7]. But those diodes do not recover the lost production. They just reduce the damage. And at the inverter level, things get stranger still: partial shading deforms the P-V curve from a single clean peak into multiple local maxima. A traditional maximum power point tracking (MPPT) algorithm, designed to find one optimum, can lock onto the wrong peak and miss a substantial fraction of available energy [5].

Figures 6 and 7 illustrate these effects. The key takeaway from Figure 6 is the disproportionate relationship: less than 5 per cent shaded area producing more than 20 to 30 per cent power loss is not a corner case; it is a predictable consequence of series string wiring under partial shading conditions [5]. Getting this right in a design requires hourly three-dimensional simulation across a full meteorological year. Simplified two-dimensional assessments simply cannot capture it. That capability requirement was a direct driver of the computational framework developed in this research.




Figure 6 Severity and impact of partial shading degradation on relative power output within a PV series string. Source: Author's elaboration.

Effect of Partial Shading on a Series String




Figure 7 Effect of partial shading on a series string and its corresponding electrical mismatch principles. Source: Author's elaboration based on Woyte et al. [5].

1.2 Building Information Modeling (BIM): Fundamentals and Evolution

BIM is one of those terms that gets used so broadly it has almost lost its meaning. In practice, it is often treated as a synonym for 3D CAD a more sophisticated way to draw buildings. That framing is understandable, but it misses what actually makes BIM useful.

The visual geometry of a building model is only a container. The value sits in the information attached to that geometry. During the development of the Revit model for the Stellantis facility examined in this research, the capability that proved most useful was not three-dimensional visualisation. Any number of software packages can do that. What mattered was being able to query a roof element directly and retrieve its structural load capacity, its obstruction clearances, its solar exposure data, and later in the workflow its simulated annual energy yield, all without opening a separate document. The model was not a picture of the building. It was a database of it.

The formal definition of BIM reflects this: the creation and management of a shared digital representation of a built asset in which geometric and non-geometric information are maintained within a common data environment throughout the project lifecycle [9]. The three-

dimensional geometry provides context. The semantic information attached to that geometry is where the analytical value lives.

For industrial PV retrofits, this matters enormously. A roof element that carries information about its load-bearing capacity, obstacle setbacks, solar irradiation potential, and operational constraints can be queried directly by downstream analytical processes. Without that integration, project teams typically maintain separate digital representations of the same building for structural analysis, energy simulation, and cost estimation. Those datasets diverge the moment the design changes, and keeping them aligned becomes a job in itself.

BIM's origins go back further than most practitioners realise. Figure 8 traces the development timeline. Conceptual precursors emerged in the 1970s the Building Description System at Carnegie Mellon and the GLIDE system that followed it were exploring object-oriented approaches to building representation at a time when CAD itself was still a novelty [9,10]. By the early 2000s, most of the principles underlying contemporary BIM platforms were already established. What has changed since is primarily the quality of the software implementation and the maturity of the workflows built around it. The underlying philosophy has not shifted: a building model's usefulness depends on the quality and accessibility of its information, not just the accuracy of its geometry.




Figure 8 Historical Building Information Modeling (BIM) development timeline. Source: Adapted from Eastman et al. [9].

1.2.1 BIM Maturity and Multidimensionality

As BIM adoption spread through the construction industry, frameworks emerged to describe what ‘doing BIM’ actually meant at different levels of organisational and technological sophistication. The Bew-Richards maturity model is the most widely cited of these [11].

The progression runs from Level 0, essentially paper drawings with a digital production tool attached, through Level 1, which introduces 3D modelling but retains independent discipline-specific workflows, to Level 2, where structured data environments allow coordinated

information exchange between project teams. Most serious BIM implementations today operate somewhere around Level 2. Level 3, in which all stakeholders interact with a single continuously updated shared model, remains more of an industry aspiration than a widespread reality [11]. Figure 9 illustrates the full spectrum.



Figure 9 The Bew-Richards BIM Maturity Model. Source: Adapted from Bew & Richards [11].

Alongside maturity levels, BIM is described through a series of ‘dimensions’ that extend what the model contains beyond three-dimensional geometry [12]. 4D adds time construction sequencing and scheduling. 5D adds cost. 6D adds sustainability and energy performance metrics. 7D adds facility management data: maintenance schedules, asset records, operational parameters. The logic is that each additional dimension should be derived from and linked to the underlying geometric model, rather than maintained as a separate dataset that periodically receives exported snapshots from it.

That principle deriving all project information from a common source rather than maintaining parallel databases is central to what this research attempts to implement. In practice, 4D and 5D workflows in particular are still often disconnected from the BIM model, receiving data through exports rather than through live connections. Closing that gap, specifically for PV retrofit projects, is one of the primary objectives of the framework developed here.




Figure 10 Dimensions of Building Information Modeling. Source: Adapted from Succar [12].

1.2.2 Level of Development (LOD) and Level of Information (LOI)

Two of the more practically useful concepts in BIM standardisation are the Level of Development (LOD) and Level of Information (LOI) frameworks, which address geometric completeness and informational completeness respectively and which are often confused with each other.

LOD describes how geometrically complete a model element is. The American Institute of Architects framework [13] runs from LOD 100 (a rough conceptual mass that indicates approximate location and scale but not much else) through LOD 200 (approximate geometry suitable for preliminary coordination) and LOD 300 (detailed geometry appropriate for engineering analysis and design development), up to LOD 400 (fabrication-level information) and LOD 500 (verified as-built condition). Broadly speaking, LOD rises as the project progresses from feasibility through design to construction.

LOI is a separate question. It asks: how much non-graphical information is attached to a model element? A simple geometric object can carry rich analytical data, while a highly detailed geometric model can contain almost none. They are independent axes. A structural beam at LOD 300 might carry complete load capacity data, deflection limits, maintenance requirements, and cost classifications. A conceptually modelled roof surface might carry only its area. Neither tells you anything about the other.

The ability to develop these two dimensions independently is genuinely useful in practice. It means that early in a project, when fabrication-level geometry would be premature and wasteful to model, the information content of the model can still be rich enough to support meaningful analytical decisions. This observation shapes the LOD 300 / LOI 300 approach adopted in this research, discussed in more detail in Section 1.3.1.

1.3 The Paradigm Shift: BIM Beyond Geometry

For most of its commercial history, Building Information Modelling (BIM) has been used primarily as a platform for design coordination, clash detection, and construction documentation. Although these functions have significantly improved project delivery, they represent only a fraction of BIM's capabilities when the model is treated as a structured database rather than merely a geometric representation. Within Autodesk Revit, the Revit API enables every model element to contain queryable and modifiable information, allowing analytical processes to interact directly with the BIM model. Consequently, geometric modifications can automatically trigger updates to energy simulations, cost estimations, and construction schedules, transforming BIM into an active decision-support platform rather than a passive documentation tool.

This capability is particularly important for photovoltaic (PV) retrofit projects, where structural constraints, roof geometry, solar resource availability, installation logistics, economic performance, and facility management are highly interdependent. Conventional workflows require each design modification to be manually propagated across multiple software environments, creating repetitive coordination tasks and increasing the likelihood of inconsistencies. The framework proposed in this research replaces this fragmented process with an API-driven workflow in which model modifications automatically initiate the required analytical updates, enabling rapid and consistent design iterations.

An important finding of this research is that increasing geometric complexity does not necessarily improve analytical performance during early-stage PV design. Annual solar simulations evaluate irradiance conditions for 8,760 hourly time steps across numerous candidate panel locations, making computational efficiency highly sensitive to model complexity. Highly detailed LOD 400 models, including fabrication components such as mounting hardware, cable conduits, and fasteners, substantially increase computational time without improving the accuracy of annual energy predictions [14]. Therefore, the proposed

methodology intentionally maintains the BIM model at LOD 300 while enriching its Level of Information (LOI). This level of geometric detail adequately represents rooftop boundaries, obstacles, structural zones, and PV module locations while maintaining efficient simulation performance. Analytical parameters—including meteorological data, hourly irradiance, module efficiency, shading losses, electrical configuration, and predicted energy production—are subsequently written back into the BIM model as shared parameters, resulting in a data-rich model that supports informed decision-making with significantly reduced computational effort. Figures 11 and 12 illustrate the complementary progression of geometric development and information enrichment within the proposed framework [13].




Figure 11 Geometric progression from conceptual representation (LOD 100) to fabrication (LOD 400). Source: Author's elaboration based on AIA specifications [13].




Figure 12 Progression of non-graphical data within the proposed BIM framework corresponding to Level of Information (LOI) tiers. Source: Author's elaboration.

Unlike conventional multidimensional BIM implementations, where scheduling, cost estimation, and energy analysis are performed in separate software environments, the proposed framework integrates these processes directly within the BIM environment using Dynamo automation, embedded Python scripting, and the Revit API. Installation schedules are automatically updated according to changes in PV layouts, reducing manual adjustments

associated with active industrial facilities. Similarly, material quantities are extracted directly from the BIM model, allowing cost estimates to evolve automatically as panel arrangements, cable lengths, and mounting requirements change. Furthermore, simulation outputs—including annual energy yield, shading losses, and electrical string characteristics—are stored as element-level parameters, providing facility managers with operational reference data throughout the building lifecycle [15,16]. Figure 13 illustrates the integrated architecture supporting continuous 4D, 5D, and 7D information exchange.




Figure 13 System architecture of the proposed API-driven multidimensional workflow, centralising 4D, 5D, and 7D data distribution. Source: Author's elaboration.

Automating PV layout generation further improves workflow efficiency. Manual placement of photovoltaic modules on large industrial roofs requires simultaneous consideration of structural limitations, maintenance access, fire-safety clearances, rooftop equipment, and inter-row shading distances, making repeated design revisions both laborious and error-prone. The proposed framework employs Dynamo, Autodesk Revit's visual programming environment, to automate this process through rule-based algorithms. Roof boundaries are extracted from the BIM model, candidate panel locations are generated, applicable design constraints are evaluated, and valid photovoltaic family instances are created automatically.

Dynamo's node-based interface enables engineers without advanced programming expertise to implement complex design logic while maintaining full access to the Revit API and DesignScript when required [17,18]. Figure 14 summarises the automated workflow from roof extraction to validated PV placement.




Figure 14 Visual programming workflow in Dynamo for extracting roof boundaries, generating spatial grids, and filtering panel coordinates against MEP obstacles. Source: Author's elaboration.

Although Dynamo efficiently manages geometric operations, computationally intensive analyses such as annual irradiance simulations require the capabilities of scientific Python libraries. Dynamo's default IronPython 2.7 environment lacks support for NumPy, making large-scale vectorised calculations impractical. To overcome this limitation, the proposed framework integrates a CPython 3.x environment through the Common Language Runtime (CLR), enabling direct use of NumPy, SciPy, and other scientific computing libraries while remaining embedded within the BIM environment [61]. Geometry extracted from Revit—including roof boundaries, rooftop obstacles, and validated panel locations—is transferred to the Python engine, where EPW weather data for the Turin region are processed to compute solar position, irradiance, shading, and annual energy production over all 8,760 hourly time steps [21]. The calculated performance indicators, including annual energy yield, shading loss factors, string allocation, and specific yield, are then written back into the BIM model using batch transaction management, which improves computational efficiency while maintaining

model integrity in the event of execution errors [19,22]. This establishes a continuous computational loop in which geometry drives simulation and simulation results immediately enrich the BIM model with queryable performance information. Figures 15 and 16 illustrate the embedded Python workflow and the integrated computational architecture.




Figure 15 Embedded Python scripting within the Dynamo visual programming environment mapping the spatial discretisation matrix. Source: Author's elaboration.




Figure 16 System architecture of the continuous loop methodology using Dynamo data routing and the embedded Python physics engine. Source: Author's elaboration

1.4 Computational Interoperability and Workflow Efficiency

The AEC industry has spent several decades and a great deal of money building sophisticated digital tools for design, analysis, and project management. The tools themselves have become genuinely impressive. The connections between them have not kept pace.

Structural analysis, energy simulation, BIM modelling, scheduling, and cost management have each evolved within their own software ecosystems. Each ecosystem does what it does well. What remains poorly solved is the handoff between them. On a simple project, that is manageable. Information moves infrequently, updates are straightforward, and the effort required to keep datasets synchronised is not overwhelming.

Industrial PV retrofits are not simple projects. Structural behaviour, roof geometry, electrical configuration, energy yield, installation logistics, and financial performance are tightly coupled. A change in one area propagates changes through all the others, often in ways that are not immediately obvious until you sit down and work through the implications manually. In conventional workflows, working through those implications manually is exactly what happens. It happens every time the design changes.

The scenario plays out like this. A structural assessment identifies a roof zone with insufficient residual capacity. The BIM model is updated. Someone then has to check whether the revised panel positions alter shading conditions. Someone else has to recalculate cable quantities. A third person updates the schedule. A fourth updates the cost estimate. None of these tasks is individually difficult, but together they represent hours of coordination work that adds no analytical value. It merely catches the rest of the workflow up to a change that was already made.

Figure 17 shows the proposed alternative: a workflow where the BIM model, simulation engine, and project management layers are continuously connected rather than periodically synchronised.




Figure 17 API-driven BIM workflow integrating Revit, Dynamo, and Python for automated design coordination, simulation, and project management updates. Source: Author's elaboration.

1.4.1 The Bottleneck of File-Based Exchange and the Iteration Penalty

Although BIM adoption has significantly improved project coordination and documentation quality, many project-management processes continue to operate outside the BIM environment. Scheduling activities are commonly developed in software such as Microsoft Project or Primavera, while cost estimation and procurement planning often rely on dedicated commercial platforms. Information exchange between these systems typically occurs through IFC exports, spreadsheet-based quantity take-offs, or other file-transfer procedures.

The limitation of this approach becomes apparent whenever a design changes.

Once information is exported from a BIM model, the connection between the source model and the receiving platform is effectively broken. What remains is a static snapshot representing the project at a particular moment in time. If modifications are introduced later, updated information must be exported again, imported into the downstream system, and reconciled with any work that has already been performed. As project complexity increases, this process becomes progressively more demanding [23,24].

Industrial photovoltaic retrofit projects are especially vulnerable to this problem. Structural assessments, operational requirements, fire-safety regulations, and evolving client priorities

frequently necessitate design revisions throughout the planning process. These changes are not exceptional events; they are a normal part of project development.

Every revision triggers a series of updates across multiple documents and software environments. Drawings, schedules, cost estimates, simulations, and procurement data can quickly become inconsistent if they are not synchronised continuously. Project teams are therefore forced to dedicate increasing amounts of time to maintaining agreement between datasets rather than improving the design itself.

Figure 18 illustrates the limitations of traditional file-based interoperability workflows and the inefficiencies that emerge when information is transferred through disconnected exchange mechanisms.




Figure 18 Limitations of traditional file-based interoperability workflows (IFC and gbXML) resulting in manual updates and workflow inefficiency. Source: Author's elaboration based on Dong et al. [23] and Venugopal et al. [24].

This research uses the term Iteration Penalty to describe what accumulates from this process. The term captures something important: the cost being incurred is not a fixed overhead. It grows with every revision, and it grows faster as more disciplines become involved. Early in a project, when there are few constraints and few stakeholders, the penalty is manageable. By the time a project is complex and the deadline is approaching, it can become the dominant factor governing how many design alternatives the team has time to evaluate.

And here is where the practical damage happens. Teams under time pressure do not explore ten design alternatives. They explore two or three, pick the best of those, and move on. The selected design may be good. But it is good relative to a very small comparison set, one determined not by engineering analysis but by the coordination overhead of generating each option.

Reducing that overhead is the central engineering motivation for the framework described in this thesis. Figure 19 shows how geometric information, simulation, and project management data remain connected throughout the workflow. There are no file exports, no manual reconciliation, and none of the lag that makes frequent iteration impractical.




Figure 19 Proposed API-driven BIM workflow integrating Revit, Dynamo, Python, and climate data within a unified computational environment. Source: Author's elaboration.

1.5 Problem Statement and Industry Bottlenecks

The primary challenge addressed in this research does not lie in photovoltaic (PV) hardware, as modules, inverters, and mounting systems are mature technologies with well-established performance characteristics. Instead, the challenge concerns the fragmented design and coordination process required to integrate these technologies into existing industrial buildings. Successful PV retrofit design requires the continuous exchange of information between multiple engineering disciplines, including BIM-based geometry, structural assessment, energy simulation, scheduling, and cost estimation. Although each discipline is supported by specialised software, communication between these environments remains largely manual, resulting in inefficient workflows, duplicated effort, and inconsistencies whenever the design evolves.

Industrial PV retrofit projects require close collaboration between structural engineers, electrical engineers, and project managers. Structural capacity must be verified according to EN 1991 [25], electrical systems must be optimised for energy production, and construction activities must be coordinated with procurement, budgeting, and operational constraints. Although BIM provides a common digital representation of the project, much of the critical analysis continues to be performed in external software. Structural assessments typically rely on finite-element analysis, energy simulations are conducted using tools such as PVsyst or EnergyPlus [26], while scheduling and cost management are maintained in separate project management platforms. Information is exchanged through formats such as IFC and gbXML, which support interoperability but often introduce data simplification, translation errors, and delays caused by manual export and import procedures [23,24]. As a result, updates generated in one discipline frequently propagate to others only after significant time lags, reducing the reliability of design decisions. Furthermore, simulation outputs—including structural utilisation, shading analysis, and predicted energy yield—are commonly stored in external reports rather than embedded within the BIM model, making it difficult to ensure that analytical results remain consistent with the latest design revision.




Figure 20 Sequence diagram illustrating the latency and rework cycles of traditional multidisciplinary communication.
Source: Author's elaboration

The consequences of fragmented information become even more significant during project management. A modification to the PV layout caused by structural constraints affects numerous downstream activities, including cable routing, mounting quantities, installation sequencing, procurement, and overall project cost. In conventional workflows these dependencies are updated independently, relying on manual coordination between multiple stakeholders. Consequently, schedules, quantity take-offs, procurement plans, and cost estimates may correspond to different design revisions, creating inconsistencies that accumulate throughout the project lifecycle. Such discrepancies frequently contribute to construction delays, budget overruns, and additional coordination effort. Figure 21 illustrates how the separation of modelling, analysis, scheduling, and cost management generates the Iteration Penalty introduced in Section 1.4.1. The methodology proposed in this research addresses this limitation by maintaining continuous synchronisation between geometric

modelling, environmental analysis, scheduling, and cost evaluation, allowing analytical and financial information to update automatically whenever the design changes.




Figure 21 Multidimensional workflow separation in conventional BIM workflows illustrating how isolated software tools result in the Iteration Penalty. Source: Author's elaboration.

Despite substantial progress in algorithmic design, energy simulation, and BIM-enabled project management, the existing literature remains largely compartmentalised. Previous studies have successfully applied Dynamo, Grasshopper, and similar computational tools to automate PV panel placement while satisfying geometric constraints; however, these approaches rarely maintain direct connections to energy simulation or project management information [20,27]. Similarly, high-fidelity solar simulation methods provide accurate irradiance modelling and performance prediction but generally operate independently of BIM environments and require repeated manual data exchange whenever the design changes [26]. Research on 4D scheduling, 5D cost estimation, and digital twins has demonstrated the value of BIM for project management, yet these studies often assume that geometric and analytical information remain continuously synchronised, an assumption that is seldom achieved in practical industrial projects.

The principal research gap therefore lies in the absence of a unified computational framework capable of integrating algorithmic PV layout generation, physically based annual energy

simulation, and multidimensional BIM-based project management within a single continuously synchronised environment. Existing workflows continue to depend on sequential file-based data exchange, resulting in repeated manual intervention, coordination delays, and inconsistent project information [23,24]. Figure 22 compares current computational approaches with the proposed methodology, demonstrating that no existing solution simultaneously provides high levels of automation, analytical integration, and multidisciplinary coordination for industrial PV retrofit applications.




Figure 22 Radar chart mapping the performance of existing computational workflows against the proposed native API-based methodology. Source: Author's elaboration.

To address this gap, the methodology developed in this thesis treats Autodesk Revit as the central information repository rather than solely as a geometric modelling platform. Through the integration of Dynamo automation, embedded Python scripting, and direct Revit API interaction, geometric design, structural information, environmental simulation, scheduling, and cost evaluation remain continuously synchronised throughout the design process. Changes introduced to the BIM model automatically propagate across analytical and management domains, reducing manual coordination, improving workflow efficiency, and enabling more informed decision-making during industrial photovoltaic retrofit projects.

1.6 Research Objectives and Methodological Framework

The framework developed in this research has one primary practical aim: to eliminate the manual data-transfer steps that currently disconnect geometric design from performance analysis and project management in industrial PV retrofit projects. Autodesk Revit was chosen as the platform because of its widespread industry adoption, its extensive and well-documented API, and its native integration with Dynamo. The adoption of CPython 3.x rather than IronPython 2.7 was not a stylistic preference but a technical necessity. Access to NumPy and the broader scientific Python ecosystem is essential for running annual irradiance calculations at the speed that supports iterative design [61].

The methodology is structured around four objectives.

The first concerns spatial analysis. Rather than estimating usable roof area through a rough exclusion percentage, the framework extracts actual roof geometry from the BIM model and subdivides it into a 0.5 m × 0.5 m grid of candidate positions. Each position is evaluated against geometric, structural, and regulatory constraints derived from the LOD 300 model [13] and EN 1991 requirements. Only positions that satisfy every constraint proceed to panel placement and the constraint evaluation is repeatable, consistent, and automatic.

The second objective is annual energy simulation. The embedded Python engine uses EPW weather data for the Turin region to calculate solar position at hourly resolution across the full meteorological year, then evaluates irradiance at each valid panel location through geometric ray tracing against the BIM model [28]. Partial shading effects are captured at module level not estimated from simplified annual factors.

The third objective is information persistence. Simulation outputs are written back into the Revit model as shared parameters on individual photovoltaic elements [18]. Annual yield, shading loss, string allocation all of it lives in the model, attached to the components it describes. The model handed over at project completion is already a reference document for operational performance monitoring.

The fourth objective is economic integration. Because technical performance data is linked directly to model elements, financial metrics (installed capacity, annual energy production, specific yield, capital cost) can be recalculated automatically whenever the design changes [15]. Cost assessment is not a post-design activity; it is part of the design loop.

1.7 Scope and Limitations of the Study

The framework developed in this research is intended as an early-stage decision-support tool for industrial photovoltaic (PV) retrofit projects. Its primary objective is to enable rapid evaluation of alternative PV layouts while maintaining sufficient analytical accuracy to support engineering decision-making. To achieve this balance, the methodology prioritises computational efficiency, automation, and comparative design assessment over high-fidelity physical modelling, thereby enabling fast iterative design without excessive computational cost.

The workflow has been developed for the industrial rooftop typologies most frequently encountered in practice, including flat, sawtooth, and parapet-edged roofs, with the San Benigno Canavese facility serving as the representative case study [29]. A spatial resolution of **0.5 m × 0.5 m** provides an effective compromise between geometric accuracy and computational performance. Although very small rooftop features may not be explicitly represented, their influence on annual PV energy production is negligible at the scale of industrial installations. Computational efficiency is further improved through spatial pre-filtering using bounding volumes before detailed geometric intersection analysis.

Environmental performance is evaluated using **EPW (TMY3)** weather data for the Turin region [30]. While Typical Meteorological Year datasets provide reliable annual energy predictions and enable comparative analyses, they do not account for short-term local effects such as transient weather events, dust accumulation, pollen deposition, or fine-scale roof temperature variations. Module temperature effects are incorporated using manufacturer-provided performance coefficients, whereas detailed airflow and thermodynamic modelling beneath the modules remain outside the scope of this study.

The proposed framework is intended to support, rather than replace, discipline-specific engineering analyses. Detailed structural verification—including wind loading assessment under EN 1991-1-4, ballast design, membrane compatibility, and seismic evaluation where applicable—must still be performed using dedicated structural engineering software [31]. Similarly, electrical engineering tasks such as cable sizing, voltage-drop calculations, inverter optimisation, protection coordination, and grid compliance verification require specialist design tools and professional engineering approval.

The principal contribution of this research is the development of a **fully integrated native BIM computational framework** that combines automated PV layout generation, embedded CPython-based solar simulation, and direct Revit API parameter write-back within a single modelling environment. By eliminating intermediate file exchanges, the framework significantly reduces the **Iteration Penalty** associated with conventional BIM–BEM workflows. It further introduces an embedded reverse ray-tracing engine with dayMask pre-processing to enable efficient large-scale solar simulation and automatically generates a 7D digital twin by embedding module-level annual energy yield, inverter assignment, and maintenance information as shared BIM parameters. Design optimisation is performed using a Multi-Criteria Decision Analysis (MCDA) approach, weighting Return on Investment (ROI) at 50%, payback period at 35%, and Capital Expenditure (CAPEX) at 15%, thereby supporting balanced technical and economic decision-making.

1.8 Research Questions

The central aim of this thesis is to determine whether an integrated BIM-centred computational workflow can improve decision-making during the early stages of industrial photovoltaic retrofit design. Current interoperability practices generally rely on IFC or gbXML exchanges to transfer information between BIM environments and external simulation tools [32]. Although widely used, these workflows introduce a discontinuity into the design process because each export represents only a snapshot of the model at a specific point in time. Any subsequent design modification requires additional exports, re-imports, and manual verification to maintain consistency across platforms.

The framework developed in this research follows a different approach. Geometric modelling, environmental simulation, and financial evaluation remain connected within a single Revit-based environment throughout the design process. By eliminating repeated file exchanges, the workflow seeks to reduce coordination effort while supporting faster and more responsive design iteration.

To evaluate the effectiveness of the proposed methodology, the study is structured around four research questions that correspond directly to the objectives outlined in Section 1.6.

RQ1: Workflow Efficiency

To what extent can an API-driven computational workflow reduce the time, effort, and coordination burden associated with conventional file-based interoperability processes? What mechanisms contribute most significantly to any observed improvements in workflow efficiency?

RQ2: Computational Accuracy

How closely do the irradiance and energy-yield results generated by the embedded Python simulation engine correspond to those produced by established external simulation platforms? Where differences occur, what factors explain the variation between approaches?

RQ3: Spatial Impact

How does the proposed grid-based spatial discretisation methodology influence photovoltaic layout generation and performance estimation when compared with conventional area-based feasibility assessment methods?

RQ4: Data Persistence and Reusability

To what extent can simulation outputs be embedded directly within BIM objects, and how effectively can these data be retrieved and reused for subsequent analysis, facility-management activities, and operational performance assessment?

Collectively, these questions examine not only the technical performance of the proposed framework but also its practical value as a decision-support environment for industrial photovoltaic retrofit projects.

1.9 Chapter Summary

This chapter has established why industrial rooftop PV retrofits are both increasingly important and persistently difficult to execute well. The economic and regulatory case for on-site solar generation has strengthened considerably over the past decade. Falling technology costs, rising carbon prices through the EU ETS [1], and growing pressure to demonstrate measurable emissions reductions have converged to make rooftop PV a mainstream capital investment for industrial facilities rather than a peripheral sustainability initiative. Piedmont's manufacturing-intensive environment exemplifies the conditions in which this argument has become most compelling.

But the chapter has also been honest about the source of the difficulty. The technology is not the problem. The process is. Industrial rooftops are geometrically complex, structurally constrained, and informationally demanding environments. The engineering disciplines required to assess them (structural, electrical, energy, project management) currently operate largely in separate software ecosystems connected by periodic file exports. Every design revision triggers a manual reconciliation effort across those ecosystems. That effort is what this thesis has defined as the Iteration Penalty, and it is the primary reason why project teams converge prematurely on suboptimal designs.

The response proposed here is an API-driven computational framework built around Autodesk Revit, Dynamo, and embedded CPython scripting. Within this environment, spatial analysis, annual performance simulation, cost evaluation, and scheduling information are not separate activities connected by file transfers. They are layers of a single continuously updated model. Design changes propagate automatically. Information stays consistent. And the engineering effort that was previously absorbed by coordination is freed up for the thing it should be doing: evaluating alternatives and making better decisions.

The framework operates deliberately at LOD 300 / LOI 300, detailed enough to support meaningful analysis and simple enough to remain computationally responsive. It is a design tool for the early stages of a project, not a replacement for the detailed engineering that follows.

The chapters that follow apply this framework to the Stellantis facility in San Benigno Canavese, evaluate its performance against the four research questions defined in Section 1.8, and assess what the results imply for BIM-centred computational practice in industrial PV retrofit design more broadly.




Figure 23 Chapter 1 conceptual overview. Source: Author's elaboration.

CHAPTER 2

Literature Review

2.1 Introduction to the Literature

Most academic writing on industrial photovoltaic design concentrates on energy yield. Planning and coordination, the part that actually decides whether a project finishes on time, gets far less attention. Yet a pattern keeps showing up in real projects: the technology works, the financial case is usually positive, and the client wants the system installed. Even so, schedules slip and budgets grow. The hardware is rarely the problem. Modules and inverters are mature, well-tested products, and their field performance can be predicted within a few percentage points [10]. What causes the trouble sits upstream, in the planning layer, where the same physical roof is described separately in structural tools, energy simulation software, cost spreadsheets, and construction schedules. None of these systems talk to each other directly.

This becomes a serious problem in retrofit work, where constraints stack on top of each other in a way that greenfield projects rarely experience. An open field is governed mostly by one thing: how much land is available. An operational industrial roof is governed by many things at once. Structural load capacity changes from bay to bay. HVAC zones cannot be touched. Fire access routes must stay clear. Noise restrictions limit when work can happen. Electrical infrastructure is already fixed in place. Production has to keep running throughout. When any single constraint forces a design change, that change ripples outward into yield calculations, cable routing, mounting quantities, installation sequencing, and cost projections. In a fragmented toolchain, someone has to catch each of these consequences by hand [32].

This is the mechanism behind what Chapter 1 calls the Iteration Penalty. Take a structural engineer who finds that a roof bay cannot support the panel density originally proposed. The layout gets revised. That revision then has to be carried, manually, into the energy simulation tool, which has to be re-run. The new yield numbers are typed into the financial model. Procurement quantities get updated separately. By the time every downstream document reflects the change, a second constraint has often already surfaced, and the whole cycle starts again. The effort needed to keep parallel documentation consistent does not grow in a straight line with the number of disciplines and design changes involved. It grows faster than that [33].

This chapter looks at how the literature has tried to deal with this problem through more integrated digital workflows. The discussion is organised around four informational dimensions: spatial coordination, construction sequencing, cost management, and lifecycle operations. These map directly onto the 3D/algorithmic, 4D, 5D, and 7D BIM dimensions that Chapter 3 brings together in one methodology. Each dimension is discussed both for what it contributes and for where it stops short of full integration. The argument that comes out of this is not simply “better tools are needed” in each domain individually. It is that the field still lacks an architecture that keeps all four dimensions computationally connected at the same time.




Figure 24 Comparison of the legacy fragmented workflow and the proposed unified API-driven loop, illustrating how continuous data integration eliminates the Iteration Penalty in industrial PV retrofit project management. Source: Author's elaboration.

2.1.1 Purpose and Scope of the Review

A literature review in an applied engineering thesis does two jobs. It maps what is already known, and it marks out what is not. The second job usually matters more for shaping the research. Cataloguing existing BIM tools and their advertised capabilities is the easy part. Pinning down exactly where those systems break down structurally is harder, but it is also the part that actually justifies developing a new methodology.

The scope here is deliberately narrow. Photovoltaic technology itself, cell physics, module degradation, inverter topology, is not covered in depth. That is not because it is unimportant. It simply does not address the problem this thesis is concerned with. The behaviour of crystalline

silicon modules under partial shading is already well understood, but those performance models exist separately from the structural representation of the roof the modules sit on. Connecting the two usually means manual work, and that work adds up into real project overhead [35].

The focus instead is on delivery and integration: how digital project management tools behave in complex construction settings comparable to industrial PV retrofits, and where they fall short. Four themes structure the chapter. Algorithmic spatial coordination looks at visual programming for automated layout generation on constrained roofs [20]. Construction logistics covers how 4D BIM links scheduling data to spatial geometry to catch conflicts before they reach the site [36]. Cost integration examines how 5D BIM connects quantity take-off to pricing, so financial estimates move with the design [37]. Lifecycle data persistence considers how 7D BIM and digital twins carry the model forward past construction into operation [38]. The chapter ends with three structured matrices that position this thesis's methodology against the existing work.

2.1.2 Review Methodology and Search Strategy

Literature was gathered from Scopus, Web of Science, ScienceDirect, and the ASCE Library, with additional targeted searches in Automation in Construction, the Journal of Information Technology in Construction, and Energy and Buildings. Together these cover a large share of the strong research at the intersection of BIM and building performance. The search window was 2003 to 2025, weighted toward post-2015 publications because that is when BIM-related tools and workflows started developing quickly. Earlier foundational papers were kept where they still set out the conceptual or methodological ground that later work builds on.

Search terms combined BIM-specific keywords (Dynamo, Revit API, IFC interoperability, 4D BIM, 5D BIM, digital twin, visual programming) with construction and energy terms (solar retrofit, BIPV, industrial facility, construction logistics, cost automation, quantity take-off). A study was included only if it showed explicit computational integration between a BIM environment and at least one downstream function, spatial, temporal, financial, or operational. Studies that only addressed isolated components, a standalone energy simulation, CPM scheduling with no spatial link, cost estimation from 2D drawings, were left out. This kept the review focused on the interoperability problem that sits at the centre of the thesis [23].

Thirty-eight studies met these criteria and are discussed in this chapter (this is the count of distinct works after removing duplicates that show up under more than one thematic column in

Table 3). Their findings feed into four matrices in Section 2.7: a master paper feature audit (Table 3), a capability-gap summary by dimension (Table 4), an integration maturity assessment (Table 5), and a cross-dimensional binary audit (Table 6). These are not meant as a simple catalogue. They are built to expose where the literature breaks apart structurally, in a way that directly feeds the methodology developed in Chapter 3.

2.2 Algorithmic Spatial Coordination and Visual Programming

Try arranging several hundred PV modules across an industrial roof by hand while also checking obstruction clearances, structural load limits, fire access corridors, and inter-row shading distances. Anyone who has done this will understand why algorithmic tools stopped being optional a while ago. The hard part is not just the volume of work. It is that no single person can reliably hold all of these interacting constraints in working memory while placing modules one at a time. Mistakes are nearly unavoidable under those conditions, and they tend to surface only after procurement or installation has already moved past the point where fixing them is cheap.

Visual Programming Languages, Dynamo in Autodesk Revit and Grasshopper in Rhino being the most common, deal with this by reversing the usual order of operations. Instead of placing geometry and checking it afterward, the designer writes the constraints as executable rules that drive the generation itself. Obstruction buffers, structural bay limits, access corridor widths, and inter-row spacing are all defined parametrically. The system then enumerates only the spatial configurations that satisfy every constraint at once, and places the elements directly in the model [20]. Compliance is not checked after the fact here. It is built into how the solution space is constructed in the first place.

A two-stage structure is typical in real implementations on industrial rooftops. First, the roof surface is broken into a fine candidate grid, 0.5 m by 0.5 m in this study, and every node is checked against a layered constraint model built from an LOD 300 BIM representation. Structural exclusions, physical obstacles, and regulatory no-go zones are applied as sequential filters, and any grid point that fails one of them is dropped. Second, placement logic runs on what is left, accounting for module orientation, inter-row spacing, and string configuration. Work that used to take significant manual drafting time is reduced to a computational step that runs in minutes and can be re-run automatically every time the upstream geometry changes [18].

This shift from manual drafting to rule-based generation has effects beyond speed. Because the constraint logic lives in code instead of in someone's judgment, it becomes auditable, reproducible, and version-controllable. Each layout can be traced back to a specific set of rules, and changing a rule produces a deterministic update to the geometry. In industrial projects, where layout decisions often have to be justified to structural engineers, fire authorities, and client's months after the fact, this traceability is not a nice-to-have. It is frequently necessary.

VPL adoption in AEC has grown a lot since roughly 2015, and the literature now covers a wide range of building types and design problems [33]. Within industrial PV retrofit workflows specifically, though, one limitation keeps showing up. Spatial optimisation is usually treated as its own isolated stage, with little or no connection to cost estimation or scheduling. The algorithm produces a geometrically valid layout, but that output is still typically exported by hand and then re-entered into financial models and planning tools. So part of the efficiency gained in the spatial computation gets eaten back up by coordination work at the interfaces.

2.2.1 Case Study: BIPV Automated Layout Optimisation in Industrial Environments

Kouhirostami and Nourian's 2022 review of BIPV workflows is one of the more thorough recent looks at how parametric design tools are actually used in industrial practice, as opposed to how they are described in methodology papers [27]. Their dataset covers forty-three implementation cases across Europe and North America, including warehouse retrofits and automotive plants, and the results push back on the assumption that VPL-based automation has already solved industrial PV layout problems.

The dominant pattern in their cases is partial effectiveness, not full reliability. Rule-based layout generation works well on geometrically simple surfaces, flat or single-level roofs with limited, uniform obstructions. As geometric complexity rises, the limits of purely rule-based logic become obvious. Sawtooth roof shapes, for example, produce shadow envelopes that shift both spatially and over time with solar position. A constraint model tuned for winter conditions at 45°N excludes a different set of viable placement zones than one tuned for summer geometry, and the gap can be large [27]. Handling this properly means coupling the geometric rule set with a solar position and irradiance engine, but most of the cases in the survey did not do this. Instead they used conservative setback heuristics, which were easier to implement but cut usable roof area and, with it, energy yield.

A second structural problem shows up after layout generation finishes. In thirty-one of the forty-three cases, the output of the spatial optimisation stayed inside the geometric modelling environment. It was never systematically linked to cost databases, procurement systems, or scheduling tools. So any change driven by a structural reassessment, client value engineering, or a regulatory revision had to be carried by hand into every downstream system. The drafting efficiency gains reported, typically forty to sixty per cent in the more mature implementations, were partly cancelled out by this manual coordination later in the workflow [27].

Castellanos, Sunter, and Kammen [3] reach a related conclusion from a different angle. Their study of rooftop PV assessment methods across several urban contexts finds that yield estimation accuracy drops sharply when spatial models skip three-dimensional obstruction geometry. Plan-view methods that apply a flat exclusion factor overestimate usable area by an average of 34% across their dataset, and the error gets worse in buildings with dense obstacles, which describes most industrial rooftops. For VPL-based workflows this has a direct consequence: candidate placement grids need to be checked against full 3D shadowing geometry, not a simplified plan projection. Coupling geometric modelling with solar simulation is not optional. It is structurally necessary.

Taken together, this body of work points to a clear gap at the spatial coordination level. Rule-based layout generation works. Its value is limited, though, unless it is tied to solar simulation and connected downstream to cost and scheduling systems.

2.3- 4D BIM: Construction Sequencing and Site Logistics

An industrial PV retrofit does not happen on an empty site. It happens on a facility that is still running production shifts, receiving deliveries, managing internal logistics, and keeping safety-critical processes going, none of which can pause for the sake of rooftop work. That changes construction sequencing in a fundamental way. What would otherwise be a fairly ordinary scheduling exercise becomes a constrained optimisation problem, with spatial, temporal, and organisational restrictions all interacting at once.

The temporal restriction is the most obvious one. Work above active production areas has to fit inside narrow windows, weekends, planned shutdowns, or shift changes. Every installation phase has to be designed around what can realistically get done in that window. Spatial constraints are usually harder to capture with conventional planning tools. A crane positioned to lift inverter equipment onto an eastern roof section might block a ground-level loading

operation at the same time. A staging area for mounting components can overlap with the delivery route. None of this shows up in a standard Gantt chart. It only becomes visible once the schedule is mapped onto a 3D model of the site over time.

This is where 4D BIM does its real work. By tying the time dimension to the spatial model, it lets the project team simulate how the site changes throughout construction. Individual elements (PV module assemblies, inverter placements, cable trays, structural reinforcement) are linked to scheduled activities. The site's planned state at any given moment can then be visualised and checked before work starts [40]. Issues that would normally surface during execution, often right when they cause a disruption, can instead be caught and resolved during planning, sometimes months ahead of time.

Moon et al. [40] show this in practice through better crane positioning and material flow planning on a large industrial project, which matters directly for rooftop PV work, where lifting operations are frequent and tightly constrained by space. Chau et al. [42] build on this by showing that linking schedule data to spatial geometry produces more realistic labour and equipment allocation than traditional CPM scheduling alone. Across the literature the outcomes reported are fairly consistent: less rework from logistical conflicts, better coordination between trades working in tight spaces at the same time, and more reliable material staging that cuts down on-site storage needs [36].

2.3.1 Case Study: Logistical Phasing in Congested, Zero-Downtime Environments

Davies and Harty's study of BIM use on a large hospital project is relevant here, not because hospitals and industrial sites are the same, but because the operational constraints line up structurally [41]. A hospital, like an industrial facility, cannot shut down for construction. Clinical work continues while building work proceeds in carefully controlled segments. In both cases, sequencing is constrained by the need to keep adjacent zones running.

Their findings show that 4D simulation let the construction programme be broken into small, highly localised work packages, each tied to a specific zone and a tightly defined time window matched to operational needs. Material deliveries were organised in phase-specific batches instead of bulk drops, which reduced blockages on access routes and cut interference with ongoing operations. One result stands out: this phasing approach, supported by 4D BIM, cut

logistical delays by roughly forty per cent compared with earlier phases of the same project that used conventional scheduling [41].

Gledson and Greenwood's wider industry survey [36] puts these findings in context. Their analysis of 4D BIM use across UK construction projects finds that practitioners consistently rate spatial-temporal conflict detection as the most valuable application of 4D tools, ahead of client communication or general progress monitoring. The conflicts that get caught are usually the ones invisible in a traditional schedule: a concrete pour planned in a spot that blocks an emergency exit route, or a crane operation scheduled for a time when the area underneath is still occupied by site staff.

Davies and Harty, Gledson and Greenwood, and Moon et al. together make a solid case for the practical value of 4D BIM. At the same time, their work points to a structural limit that is not always stated outright. In most implementations, 4D simulation is built in a separate piece of software from the main BIM authoring platform. Tools like Navisworks and Synchro typically import IFC or NWC exports from Revit, and schedule data is linked to build the 4D model independently of the live design environment [39]. Once that separation exists, model synchronisation stops being continuous. Any design change requires a re-export, re-linking of schedule relationships, and rebuilding the simulation.

Where the design is stable before construction starts, this workflow is fine. In industrial PV retrofits, that assumption rarely holds. Structural reassessments, operational constraints, procurement lead times, and sequencing adjustments routinely force design changes well after construction has begun. Under those conditions, the gap between design and 4D simulation becomes a real limitation. Botton et al. [39] make this point directly in their study of logistics simulation for complex projects, noting that 4D models lose their usefulness fast when they cannot be updated in near real time alongside the evolving design. The implication is simple: a 4D model that depends on periodic file exports will, in an iterative industrial setting, always lag behind the actual state of the project.

2.4 -5D BIM: Cost Integration and Live Financial Estimation

Cost management in industrial PV retrofits sits in an awkward spot. Capital expenditure is tied directly to panel count and layout, but those numbers rarely settle until fairly late in design. At the same time, procurement decisions for long-lead items like inverters and specialised

mounting systems cannot wait for full design certainty. In practice, financial commitments often get made while the geometry underneath is still moving.

The usual response is a separate cost model kept in a dedicated financial tool, updated by hand every time the BIM model changes. In theory this keeps design and estimation cleanly separated. In practice it mostly means the cost model is a few iterations behind reality at any given moment [37].

5D BIM tries to close this gap by linking model objects, PV panels, inverters, racking, cable trays, directly to cost databases. Quantities get pulled automatically from the geometry and mapped to unit rates, producing cost outputs that move with the model. Smith [37] frames this as a change in what the cost manager actually does day to day. Instead of measuring off drawings, the job becomes more about querying the model directly. Mitchell [46] takes this further, arguing that once cost is continuously visible, design iteration behaves differently, because every geometry change now carries an immediate financial consequence.

The literature is more cautious in practice. Lu et al. [43] and Kim et al. [45] both find that 5D BIM usually does not run as a continuous live system. It tends to operate at fixed checkpoints in the design process instead. The reason is fairly structural. Most implementations rely on plugins or IFC exports that generate a point-in-time quantity take-off. Once the design changes, which happens often in PV retrofit work, the whole process has to be repeated: quantities exported again, the cost model updated again, consistency rebuilt by hand. The “live” part ends up being more of a repeated snapshot than a continuous connection.

Vigneault et al. [44] add another piece to this, one that matters a lot for PV systems specifically. Their work on lifecycle cost integration shows that operational costs are usually left out of the BIM model entirely during construction. Maintenance cycles, inverter replacement, production losses during downtime, these tend to sit outside the model altogether. So investment decisions often get made from a CAPEX-only view, even though the real financial picture is wider than that. Writing operational outputs back into the model as 7D parameters, the approach taken in Chapter 3, closes that gap more structurally than typical 5D setups allow.

2.5- 7D BIM: Lifecycle Data and Digital Twin Integration

The 7D dimension covers what happens after handover, although in practice handover is usually where active use of the BIM model starts to fade. The model still exists. It is technically complete. But it gets passed to facility management and then slowly drifts away from day-to-

day operational decisions. The mismatch shows up quickly: the model is rich in geometry and specification detail, but it is not really built to match how the building is actually operated.

This is very visible in industrial PV systems. An underperforming inverter does not show up as a clear state inside the BIM model. It shows up in SCADA or an energy monitoring system instead. Design intent sits in one place, operational reality in another. When expected and actual yield diverge, there is no shared structure holding both perspectives together. Boje et al. [38] call this a semantic continuity problem in digital twins. The hard part is not building the model. It is keeping information consistent across the building's life.

Digital twin research has gone fairly far on this point. Khajavi et al. [48], Lu et al. [47], and Cavka et al. [15] all contribute to the idea of a continuously updated digital counterpart, where sensor data feeds back into the model and refreshes its assumptions over time. Aerospace and large infrastructure are already well ahead here. Buildings are catching up. Lu et al.'s West Cambridge Campus study [47], for example, shows that IoT integration can support real-time updates to model parameters. Even so, one limitation keeps showing up across this work: the BIM model used during construction is usually not the same model used during operation. A second model gets built instead, which means duplicated effort and a break in continuity.

Gerrish et al. [50] and Parn et al. [51] show this clearly for energy performance. BIM models rarely get used directly inside facility management systems, not because the data is useless, but because it is not structured in a way FM platforms can consume easily. So data gets re-entered by hand, even though it already exists somewhere in the model. That format mismatch is where a lot of value gets lost.

The approach in Chapter 3 tries to address this earlier in the process. Instead of exporting data after the fact, simulation outputs, annual yield, shading losses, string-level performance, are written back into the Revit model as shared parameters linked to PV elements. That turns the BIM model into a 7D container during design, not after it. For facility management, this means performance expectations are already sitting in the same environment used to design the system, with no separate reconstruction step needed.

2.6 The Integrated Research Gap: Synthesis

Look across all four BIM dimensions and the pattern holds up consistently. Each domain is mature on its own. Spatial optimisation works. 4D sequencing works. 5D cost mapping works. Digital twin concepts are well developed. The problem is not that any of these pieces are

missing or underdeveloped. It is that they do not stay connected once you move past a single tool.

What exists in practice is a set of isolated systems? A 3D layout tool produces geometry but does not talk directly to the 4D scheduling environment. The 4D model is separate again from cost estimation, and the cost model is not naturally tied to lifecycle data structures. Information moves between them through exports, IFC, gbXML, NWC, CSV. These work fine at the moment of export, but drift starts right after. Each tool ends up working on a frozen snapshot of the project, and the person running the project becomes responsible for keeping everything consistent by hand.

This is not a new observation. Dong et al. [23] flagged IFC interoperability limits back in 2007. Kamel and Memari [32] described the BIM-to-BEM disconnect. Kim et al. [45] showed that plugin-based 5D systems hit a ceiling where integration stays partial. What has changed more recently is not the diagnosis but the possibility of fixing it. The Revit API is now mature enough to support real bidirectional control of model data. Dynamo with embedded CPython opens access to a full scientific computing stack inside BIM. At the same time, case study after case study shows that partial integration does not remove the Iteration Penalty, because the parts that stay unconnected still need manual coordination.

So the gap is fairly specific. No existing workflow connects algorithmic PV layout generation, solar simulation, automated 5D cost extraction, and 7D lifecycle data embedding inside one continuously synchronised Revit-based environment for industrial PV retrofits. The individual pieces exist. They have just never been brought together in a closed loop.

2.7 Literature Matrices and Summary

The tables below synthesise the thirty-eight studies reviewed in this chapter. Table 3 is the master feature audit, mapping every study against the five integration dimensions covered here. Table 4 summarises capability and gap by thematic cluster. Table 5 assesses integration maturity and iteration support. Table 6 is a binary cross-dimensional audit of the twelve most cited studies alongside the proposed methodology. Together, these tables make the structural gap visible in a form that directly motivates the design choices made in Chapter 3.

In Tables 3 and 6: ✓ = capability present and demonstrated; ✗ = absent; ● = partial or emerging; Part = partial implementation noted in the study.

Table 3 Master paper feature audit: all reviewed studies mapped against integration dimensions.

Ref	Author(s) & Year	Primary Domain	3D / Spatial	4D Sched.	5D Cost	7D / FM	API / Live Link	Key Limitation
[10]	Sacks et al., 2018	BIM Handbook / Foundations	✓	✗	✗	✗	✗	Reference text; no implementation
[32]	Kamel & Memari, 2019	BIM–energy simulation review	✓	✗	✗	✗	✗	Survey only; no integrated workflow
[33]	Cui et al., 2021	Visual programming in AEC	✓	✗	✗	✗	Part	VPL review; no cost/schedule link
[34]	Lu et al., 2013	BIM & construction practice change	✓	✗	✗	✗	✗	Conceptual; pre-API era
[35]	Sanjuán-Micó et al., 2020	BIM-to-BEM workflow	✓	✗	✗	✗	✗	File-based export; no live sync
[20]	Kensek, 2015	VPL: energy & shading case studies	✓	✗	✗	✗	Part	Geometry only; no cost/FM output
[18]	Preidel & Borrmann, 2016	Code compliance via VPL	✓	✗	✗	✗	Part	Compliance checks; no simulation link
[27]	Kouhirostami & Nourian, 2022	BIPV automated layout review	✓	✗	✗	✗	✗	Spatial only; 31/43 cases unlinked to cost
[19]	Asl et al., 2015	BIM-based performance optimisation	✓	✗	✗	✗	Part	Single-dimension optimisation; no schedule
[3]	Castellanos et al., 2017	Rooftop PV spatial assessment	✓	✗	✗	✗	✗	2D exclusion factor; no 3D shadows
[36]	Gledson & Greenwood, 2016	4D BIM adoption survey (UK)	✓	✓	✗	✗	✗	Export-based; no live design link
[39]	Boton et al., 2015	4D BIM logistics simulation	✓	✓	✗	✗	✗	Simulation lags design changes
[40]	Moon et al., 2014	4D CAD: crane & logistics	✓	✓	✗	✗	✗	CAD-based; no BIM data model
[41]	Davies & Harty, 2013	Site BIM: hospital case study	✓	✓	✗	✗	✗	IFC export dependency; no cost link
[42]	Chau et al., 2004	4D resource utilisation system	✓	✓	✗	✗	✗	Legacy system; no BIM integration
[37]	Smith, 2014	BIM & the 5D cost manager	✓	✗	✓	✗	✗	Manual QTO export; no live update

Ref	Author(s) & Year	Primary Domain	3D / Spatial	4D Sched.	5D Cost	7D / FM	API / Live Link	Key Limitation
[43]	Lu et al., 2016	5D BIM cost management	✓	✗	✓	✗	✗	IFC-based; snapshot not live
[37]	Smith, 2014	BIM 5D cost certainty	✓	✗	✓	✗	✗	Conceptual; no API link shown
[44]	Vigneault et al., 2020	Life-cycle costing in BIM	✓	✗	✓	✗	✗	LCC framework; no real-time update
[45]	Kim et al., 2020	Integrated 5D cost & schedule	✓	✓	✓	✗	Part	Partial API; manual reconciliation remains
[46]	Mitchell, 2012	5D BIM: cost certainty & buildings	✓	✗	✓	✗	✗	Industry guide; no computational demo
[38]	Boje et al., 2020 (a)	Semantic Construction Digital Twin	✓	✗	✗	✓	✗	Conceptual framework; no live BIM link
[38]	Boje et al., 2020 (b)	Digital twin directions for research	✓	✗	✗	✓	✗	Research agenda; not implemented
[15]	Cavka et al., 2017	Owner information requirements BIM	✓	✗	✗	✓	✗	Information model; no API automation
[47]	Lu et al., 2020	Digital Twin at building/city level	✓	✗	✗	✓	Part	IoT linked; construction phase not covered
[48]	Khajavi et al., 2019	Digital twin vision & boundaries	✓	✗	✗	✓	✗	Conceptual; no retrofit context
[49]	Andriamamonjy et al., 2018	Automated BIM→BEM generation	✓	✗	✗	✗	Part	Automated but file-export pipeline
[50]	Gerrish et al., 2017	BIM for existing building EPC	✓	✗	✗	✓	✗	Retrofit focused; no live cost/schedule
[51]	Parn et al., 2017	BIM trajectory in FM	✓	✗	✗	✓	✗	FM review; construction not integrated
[56]	Jalaei & Jrade, 2015	BIM + LEED at conceptual stage	✓	✗	✗	✗	✗	LEED metrics only; no solar/industrial
[23]	Dong et al., 2007	IFC vs gbXML comparative study	✓	✗	✗	✗	✗	Format comparison; highlights data loss

Ref	Author(s) & Year	Primary Domain	3D / Spatial	4D Sched.	5D Cost	7D / FM	API / Live Link	Key Limitation
[1]	EDGAR / Climate DB, 2023	Industrial GHG emissions data	X	X	X	X	X	Data source; no workflow content
[2]	IRENA, 2023	Renewable capacity projections	X	X	X	X	X	Data source; no BIM content
[3]	Castellanos et al., 2017 (spatial)	Rooftop obstacle classification	✓	X	X	X	X	Obstacle taxonomy; no BIM integration
[4]	EN 1991 (Eurocode)	Structural load standards	X	X	X	X	X	Normative standard; no workflow
[52]	Bidram et al., 2012	Bypass diode modelling	X	X	X	X	X	Module level; no BIM context
[53]	Silvestre & Castañer, 2007	String shading & bypass diodes	X	X	X	X	X	Electrical model; no BIM link
[54]	Petrone et al., 2007	MPPT under partial shading	X	X	X	X	X	Control algorithm; no BIM context
[5]	Woyte et al., 2003	Partial shading of PV arrays	X	X	X	X	X	Foundational shading study; pre-BIM
[9]	Eastman et al., 2011	BIM handbook (1st/2nd ed.)	✓	X	X	X	X	Foundational text; no API workflows
[55]	Sampaio et al., 2016	LOD vs simulation accuracy	✓	X	X	X	X	LOD study; no solar or cost link

Table 4 Capability and gap summary by thematic cluster.

Stream / Cluster	3D Spatial	4D Schedule	5D Cost	7D / FM	Structural Gap
Algorithmic VPL / spatial (§2.2) [18,19,20,27,33,35]	✓ Strong	X None	X None	X None	Layouts generated but not connected to cost, schedule, or FM data
4D scheduling & logistics (§2.3) [36,39,40,41,42]	✓ Strong	✓ Strong	X None	X None	4D model updated via IFC export; lags design; no cost propagation
5D cost integration (§2.4) [37,43,44,45,46]	✓ Strong	Partial	✓ Strong	X None	QTO export is manual; cost model freezes at export moment
7D / digital twin / FM (§2.5) [15,38,47,48,49,50,51]	✓ Strong	X None	X None	✓ Strong	FM linked post-handover; no live construction-phase update loop

Stream / Cluster	3D Spatial	4D Schedule	5D Cost	7D / FM	Structural Gap
BIM foundations (§2.1) [9,10,23,32,34]	✓ Conceptual	✗ None	✗ None	✗ None	Theoretical frameworks; interoperability treated as file problem
THIS STUDY — API-driven [Proposed methodology, Ch. 3]	✓ LOD 300	✓ Embedded	✓ Embedded	✓ Embedded	All four dimensions within a single live Revit/Python environment

Table 5 Integration maturity assessment of key studies.

Study	Interop. Method	Update Mechanism	Dimensions Linked	Maturity Level	Iteration Support	Best-fit Use Case
Kouhirostami & Nourian [27]	None (geometry only)	Manual re-export	3D only	Low	✗	Layout feasibility
Preidel & Borrmann [18]	VPL → model	Re-run script	3D only	Low–Med	Partial	Compliance checking
Gledson & Greenwood [36]	IFC export	Periodic re-sync	3D + 4D	Medium	Partial	Site logistics planning
Moon et al. [40]	CAD file export	Manual re-link	3D + 4D	Low–Med	✗	Crane & delivery sequencing
Boton et al. [39]	IFC export	Manual re-sync	3D + 4D	Medium	✗	Construction logistics
Kim et al. [45]	Partial API	Semi-automatic	3D+4D+5D	Med–High	Partial	Cost & schedule integration
Boje et al. [38]	Semantic model	Conceptual only	3D + 7D	Concept.	✗	Digital twin theory
Lu et al. [47]	IoT + BIM	Real-time sensors	3D + 7D	High (ops)	✗	Operational monitoring
Andriamamonjy et al. [49]	Script → IFC → BEM	Automated but file-based	3D only	Medium	✗	Energy model generation
THIS STUDY	Native Revit API	Fully automatic	3D+4D+5D+7D	High	✓	Industrial PV retrofit design

Table 6 Binary cross-dimensional audit. ✓ = present; ✗ = absent; ◐ = partial. Source: Author's elaboration.

Study	3D BIM Model	Solar Sim.	Structural Check	4D Schedule	5D Cost	7D FM	API / Live	Industrial Context	Fully Integrated
Kouhirostami & Nourian [27]	✓	✗	✗	✗	✗	✗	✗	✓	✗
Kensek [20]	✓	✓	✗	✗	✗	✗	◐	✗	✗
Castellanos et al. [3]	✓	✓	✗	✗	✗	✗	✗	✗	✗
Gledson & Greenwood [36]	✓	✗	✗	✓	✗	✗	✗	✓	✗

Study	3D BIM Model	Solar Sim.	Structural Check	4D Schedule	5D Cost	7D FM	API / Live	Industrial Context	Fully Integrated
Moon et al. [40]	✓	✗	✗	✓	✗	✗	✗	✓	✗
Boton et al. [39]	✓	✗	✗	✓	✗	✗	✗	✓	✗
Kim et al. [45]	✓	✗	✗	✓	✓	✗	🟡	✓	✗
Lu et al. [47]	✓	✗	✗	✗	✗	✓	🟡	✗	✗
Boje et al. [38]	✓	✗	✗	✗	✗	✓	✗	✗	✗
Andriamamonjy et al. [49]	✓	✓	✗	✗	✗	✗	🟡	✗	✗
Gerrish et al. [50]	✓	✓	✗	✗	✗	✓	✗	✓	✗
THIS STUDY	✓	✓	✓	✓	✓	✓	✓	✓	✓

2.8 Chapter Summary

This chapter reviewed thirty-eight studies across four BIM dimensions: spatial coordination, 4D scheduling, 5D cost integration, and 7D lifecycle management. One thing stands out consistently across the literature. Each dimension on its own has reached a fairly mature state, with established methods and a wide range of applications. The trouble starts once these dimensions are expected to run together inside one synchronised computational environment. That level of integration does not really exist yet in the work reviewed here, at least not in a stable, continuous form. The tables in Section 2.7 lay this separation out explicitly.

The same pattern repeats across all four dimensions. Interoperability is usually achieved through file exchange rather than direct programmatic coupling between systems. Formats like IFC and similar exports serve as transfer mechanisms, but they only hold consistency at the exact moment the file is generated. Once the design moves on, that consistency breaks down. Downstream outputs drift away from the current state of the model and need manual correction. In effect, this brings back a large share of the coordination work that digital workflows were supposed to eliminate in the first place.

This is what the thesis calls the Iteration Penalty. It is not one isolated inefficiency. It builds up gradually, every time a design change has to cross from one disconnected tool to another. On a simpler project this might stay manageable. In an industrial PV retrofit it becomes much more visible, because structural constraints, scheduling requirements, and procurement dependencies are all interacting in parallel.

Chapter 3 responds to this directly. Rather than treating file exchange as a necessary layer between tools, it treats file exchange as the main source of the fragmentation in the first place. The proposed system keeps spatial, temporal, cost, and lifecycle representations inside a single Revit-based environment, connected through the Revit API and a CPython simulation engine. Within that setup, design changes propagate automatically across all four dimensions, with no manual export step and no re-entry of data. How that architecture is actually built is the subject of Chapter 3.

CHAPTER 3

Methodology and Computational Framework

3.1 Introduction and System Architecture

Anyone who has actually managed a construction project knows that the initial plan rarely survives contact with reality. A structural survey turns up conditions that are not on the drawings. A tender comes back over budget. A specified product becomes unavailable right at procurement. None of this is unusual. It is the normal rhythm of industrial projects. The real question is not whether change happens, but how well the planning environment absorbs it.

What has changed in recent years is the sheer density of digital tools surrounding a design team. The BIM model lives in one platform. Energy simulation runs in another. Cost estimation sits somewhere else, and scheduling is managed in yet another system. Each tool works well on its own, but the connections between them are weak. When a change comes in, the response is still mostly manual: export the model, re-enter the quantities, rerun the cost numbers, check whether the schedule still holds. Every one of these transfers carries a small risk of losing data or introducing an inconsistency. Add up enough design iterations, and these small risks compound into what this thesis calls the Iteration Penalty.

This chapter starts from a different assumption. Geometry, simulation, and cost evaluation should not be treated as separate systems passing static files back and forth. They work better as parts of one computational environment that stays synchronised as the design evolves. When the layout changes, downstream outputs should update on their own. When a financial scenario is modified, it should still be built on the same geometric and environmental data, not an outdated snapshot. The sections that follow describe how such an environment was built and applied to the Stellantis facility at San Benigno Canavese.

3.1.1 The Tool Ecosystem: A Tripartite Architecture

The system is built around three core tools, each with a specific role. They communicate through direct API calls rather than file exchange. Table 7 summarises their functions, their responsibilities, and the reasoning behind their integration.

Table 7 Summary of tools, their roles and design rationale within the tripartite computational architecture. Source: Author's elaboration.

Tool / Source	Role in Framework	Primary Function	Key Capability	Design Rationale
Autodesk Revit [61]	Central BIM database	Stores all building geometry, receives computed results as shared parameters	Roof LOD 300 model; IFC-compatible output; authoritative record for all data	Single source of truth; no data outside the model
Dynamo [17]	Data router / geom. processor	Extracts surfaces, builds sensor grids, filters constraints, routes data to Python	UV subdivision; vector filtering; FamilyInstance placement; API write-back	Visual scripting of geometry logic without manual scripting
CPython 3.x	Computational engine	Ray-tracing, irradiance integration, 135-scenario MCDA, yield normalisation	NumPy vectorisation; EPW parsing; MCDA ranking; CAPEX calculation	Full scientific stack (NumPy, SciPy) — impossible with IronPython 2.7
EPW Dataset [30]	Meteorological input	8,760 hourly climate records for Turin-Caselle; drives solar position and irradiance	GHI, DNI, DHI, T_amb at 1-hour resolution; TMY3 standard	Actual seasonal/diurnal variation rather than monthly averages
Prezziario Regionale Piemonte [57]	Regional cost database	Local labour rates, equipment costs, and infrastructure pricing for all CAPEX outputs	Unit rates for panels, inverters, racking, cable trays, labour	Grounded in actual Piedmont market conditions

Within this architecture, Autodesk Revit is not simply a drafting tool. It works as a structured relational database. All physical building components, including roof geometry, parapet structures, and HVAC obstruction volumes, are defined and managed inside it. It also acts as the final repository for computed outputs. Irradiance results, financial indicators, and 7D facility management attributes are written back into the model as IFC-compatible shared

parameters, instead of being stored somewhere outside it. In this sense, the BIM model is the authoritative record. Analytical data does not exist outside it.

Dynamo sits as the intermediary layer, handling geometry-intensive operations that would otherwise need extensive manual scripting. This includes extracting roof surfaces from the Revit model, subdividing them into spatial sampling grids, filtering viable panel locations against obstruction constraints, and structuring the data for transfer to the Python engine. It also manages two-way communication between Revit and the computational core, so processed outputs are correctly written back into the BIM environment [17]. Dynamo does not do heavy computation. It makes sure data moves correctly between the system components.

CPython 3.x is embedded inside the Dynamo workflow to handle the numerically intensive tasks that visual scripting cannot manage. This includes processing 8,760 hourly meteorological records, running ray-tracing across 8,460 spatial points, and evaluating 135 financial configurations against a regional cost database [19]. Integration with the Revit API means results are written back into the BIM model as soon as they are computed. There is no exporting step, no manual reconciliation. The feedback loop between analysis and design stays continuous. Choosing CPython over Dynamo's native IronPython 2.7 environment was a deliberate decision. IronPython does not support NumPy, and without vectorised array operations, the ray-tracing workload would be computationally impractical.




Figure 25 Data exchange architecture showing bidirectional communication between Autodesk Revit, Dynamo, and an embedded Python engine integrated with external climate and regional cost data. Source: Author's elaboration.

3.1.2 The Continuous Feedback Loop

The practical effect of this architecture shows up most clearly when design changes happen. If a structural engineer adds a new exclusion zone around a rooftop penetration, or the baseline energy demand is revised, there is no need to export updated IFC files or restart the workflow on a separate platform. The system re-runs inside the same environment. It regenerates the solar irradiation map, updates the Bill of Quantities, recalculates CAPEX, and adjusts the 4D

installation schedule directly from the modified geometry. Decisions are then based on live outputs, not on static spreadsheets that may already be out of date.

Table 8 maps the full eleven-stage workflow, from BIM baseline construction through to cross-platform validation, and lists the tool used at each stage, the key output, the computational method, and why that step matters within the overall framework.

Table 8 Complete workflow stage summary showing tools, outputs, methods, and significance for each of the eleven computational stages. Source: Author's elaboration.

#	Stage	Tool	Key Output	Method / Parameter	Significance
1	BIM Baseline Build	Revit	LOD 300 roof model; validated total area 26,372.32 m ²	FootPrintRoof + ExtrusionRoof; parapet, HVAC volumes, expansion joints	Authoritative geometry; no subsequent measurement divergence
2	Surface Filtering	Dynamo	Upward-facing faces only (facesUp)	Normal vector dot product vs Z-axis > 0.5 threshold	Eliminates parapets, soffits, edge returns automatically
3	Sensor Grid	Dynamo	8,460 spatial analysis points (PtsFlat)	UV domain subdivision 5×5 per surface; proportional to face dimensions	Location-specific irradiance — no area-averaged approximation
4	Packing Factor	Dynamo / Python	Net constructible area 23,207.64 m ² (packing factor 0.88)	12 % excluded: maintenance walkways, inverter zones, cable routing	Realistic installation density rather than geometric ceiling
5	Day Mask	Python (CPython)	Daylight hours only: 8,760 → 4,388 hours retained	Filters hours where GHI = DNI = DHI = 0 (all nighttime)	~50 % reduction in ray-casting load; zero physical accuracy loss
6	Ray-Tracing	Python (CPython)	Full annual irradiance per sensor point; peak 1,891 kWh/m ² /year	Reverse ray-cast: sensor → sun; albedo $\rho = 0.2$; $dt = 1$ h	Resolves actual shadow geometry from LOD 300 model

7	Yield Normalisation	Python	Site-wide weighted average 169.03 kWh/m ² /year	$Y_{\text{face}} = (\text{kWh/m}^2_{\text{face}}) \times \text{PR}(0.80) \times \text{Eff}(0.20)$, where kWh/m ² _face is the arithmetic mean of the 36 per-point irradiance accumulations within each surface — Converts raw irradiance to deployable electricity	Converts raw irradiance to deployable electricity
8	MCDAs (135 scenarios)	Python	Rank 1: 450 W panels, 38×100 kW string inverters, no BESS	String vs. central inverters; panel Wp; BESS inclusion; Prezziario pricing	Evidence-based selection across full configuration space
9	Module Placement	Dynamo → Revit	8,390 modules placed at validated sensor-grid coordinates	FamilyInstance.ByPoint from PtsFlat; heat-map cross-referenced	Layout generated from analysis, not designed then checked
10	7D Semantic Embedding	Dynamo API	Annual yield, inverter ID, maintenance tag on each BIM element	Element.SetParameterByName → shared IFC parameters	FM-ready model: design intent queryable against live operational data
11	Cross-Platform Validation	Grasshopper / Ladybug [58]	Spatial agreement confirmed; total within engineering tolerance	Same EPW + Turin geometry in Ladybug Sky Matrix; mesh vs. point-cloud	Physical reliability confirmed against established industry tool

3.2 Case Study Context: The Stellantis San Benigno Canavese Facility

Testing a new computational method on a simplified, purpose-built case study would have narrowed the claims that could be made about its practical relevance. The Stellantis automotive manufacturing facility at San Benigno Canavese was chosen for the opposite reason: it resists simplification. It is a fully operational industrial site, with dense rooftop installations, complicated geometry, restricted access zones, and a large continuous energy demand—in short, the type of environment where fragmented conventional workflows tend to fail most visibly. **This site also serves as a collaborative research hub, where several other students and researchers are simultaneously investigating complementary aspects of its**

operational efficiency and structural optimization. If the integrated approach holds up here, its use in more routine contexts becomes a much easier case to make.

3.2.1 Geographical and Environmental Profile

San Benigno Canavese sits in the Piedmont plain north of Turin, close to where the land starts rising into the Alpine foothills. The local climate is temperate continental. Summers bring relatively high solar irradiance. Winters are dominated by persistent cloud cover, valley fog, and reduced atmospheric transparency. Both extremes matter for the simulation. High summer irradiance sets the peak generation potential, but winter conditions, particularly diffuse irradiance and shading, are just as important when estimating annual yield without introducing a systematic overestimate.

The simulation uses the EnergyPlus Weather dataset for Turin-Caselle airport (ITA_Torino-Caselle.160590_IGDG.epw) [30], which provides hourly records of Global Horizontal Irradiance (GHI), Direct Normal Irradiance (DNI), Diffuse Horizontal Irradiance (DHI), and ambient air temperature across a full Typical Meteorological Year, 8,760 data points in total. A measured TMY dataset captures the real seasonal and diurnal variation at this location. Monthly averages would have introduced a systematic bias into the yield estimates, the kind of shortcut that tends to cause underestimation in less careful approaches [64]. The Italian national photovoltaic database from PVGIS [59] was used as a secondary cross-reference, to check that the TMY3 figures were consistent with long-run regional solar resource data.



Figure 26 Aerial context of the Stellantis manufacturing facility in San Benigno Canavese, Piedmont.




Figure 27 Master plan and architectural site layout of the industrial building complex used as the case study for the rooftop PV retrofit project.

3.2.2 Architectural Complexity and Existing Structural Topology

An active automotive plant presents a fundamentally different spatial challenge than open-field photovoltaic installations or typical commercial rooftops. The Stellantis facility cannot be treated as an empty surface waiting to be filled. It is already occupied, and every rooftop element exists to serve some operational purpose.

Aerial surveys and site documentation show a multi-level roofscape made up of several distinct structural zones. Sawtooth roof profiles, originally built decades ago to bring daylight into deep production halls, sit alongside flat roof sections and areas of varying slope that reflect different construction phases and structural systems. On top of this geometry sits a dense layer of existing infrastructure: large HVAC units, ventilation stacks, extraction systems, full-span structural expansion joints, and safety corridors that must stay clear at all times. These constraints are not negotiable in practice. They are either structurally essential or operationally required, and the photovoltaic layout has to work around them rather than the other way around.

The structural implications go beyond geometry. Adding more than 8,000 photovoltaic modules, plus their mounting systems and electrical infrastructure, adds a substantial dead load to a steel framework that was never designed with this retrofit in mind. The design therefore

had to balance energy objectives against structural capacity limits continuously. This constraint was treated as a hard boundary throughout the optimisation, not as something checked only at the end.



Figure 28 Aerial satellite view of the industrial plant site, highlighting the 25,000 m² rooftop footprint evaluated for the photovoltaic retrofit installation.




Figure 29 Breakdown of the federated BIM model into its constituent electrical, mechanical, and architectural sub-models for coordinated engineering analysis. Source: *BIM-to-BEM Framework for Energy Retrofit in Industrial Buildings: From Simulation Scenarios to Decision Support Dashboards*

3.2.3 Digitisation of the BIM Baseline

Before any spatial computation could happen, the physical facility had to be turned into a reliable digital model. A detailed baseline BIM model of the site was built in Autodesk Revit, representing the structural steel framework at LOD 300, a level of detail high enough for clash detection and coordination, but not so high that it would slow iterative simulation down to impractical levels [13]. Roof geometry was modelled as discrete elements, separating FootPrintRoof and ExtrusionRoof categories to capture variations in slope, construction logic, and structural form accurately.

At this stage the model contained no photovoltaic components or retrofit elements. It represented only the existing built condition. The total roof area was pulled directly from Revit Schedules before any external processing, giving a verified baseline value of 26,372.32 m² that the rest of the computations could rely on. When Dynamo-based geometry extraction was applied later, it referenced the same underlying surfaces already defined and validated inside the BIM environment, rather than relying on exported or duplicated geometry that could drift from the source model. In practical terms, this is what is meant by a single source of truth: every derived quantity stays traceable to the same geometric entities, and any change to those entities propagates through the entire analytical chain automatically.




Figure 30 Extracted roof category from Revit LoD300 Model

3.2.4 Energetic Baseline and Managerial Directives

Direct access to the facility's utility billing data was not available. This is a fairly common limitation in industrial settings, where confidentiality policies restrict access to operational energy records. The analysis therefore proceeded without primary billing data, and the annual energy baseline was reconstructed using two independent approaches, which were then compared to see whether they converged in a meaningful way.

The first approach was top-down benchmarking. Published Specific Energy Consumption (SEC) indicators for European automotive manufacturing facilities, particularly those involved in assembly, metal processing, and paint shop operations, were used as the reference. In this sector, SEC values typically fall between 400 and 500 kWh/m²/year [60]. Applying this range to the confirmed operational footprint of more than 30,000 m² gives a first-order estimate of roughly 15,000,000 kWh per year.

In parallel, a bottom-up load breakdown was built. Instead of starting from area-based benchmarks, this approach reconstructs consumption from typical subsystem-level behaviour in automotive plants. HVAC and industrial ventilation make up the largest share, around 40% of total demand, because of continuous operation across large production halls. Robotics and manufacturing equipment, including welding lines, stamping presses, and conveyor systems, contribute roughly 25%. Compressed air systems for pneumatic tooling account for about 15%, and lighting plus auxiliary services make up the remaining 20%. Assuming a realistic operating regime of two to three shifts per day, around 4,000 to 6,000 operating hours a year, the bottom-up estimate lines up closely with the top-down result. This convergence between two different estimation paths gives a reasonable basis for fixing annual electricity consumption at 15 GWh, which is then used as the primary design target for the retrofit. Because the internal document is confidential [29], this baseline carries an estimated uncertainty margin of $\pm 5\%$. Still, the agreement between the top-down and bottom-up methods gives a fairly solid working parameter.

All CAPEX and ROI calculations were linked directly to the Prezziario Regionale Piemonte [57], the official construction cost database for the Piedmont region, so that the financial modelling reflects local market conditions rather than generic international assumptions. In practice this matters: differences between regional and national cost baselines can shift estimated payback periods by several months for a project of this scale.

3.3 Algorithmic Spatial Discretisation: The Geometric Engine

Solar irradiance across a complex industrial roof is rarely uniform, and treating it as if it were leads to distortions in predicted performance. A single averaged value cannot capture the difference between a flat roof section in full sun from 09:00 to 16:00 and a sawtooth face that stays partly shaded for several hours in the morning because of an adjacent parapet. To capture these spatial variations properly, the roof has to be broken down into a large number of independent evaluation points rather than treated as one continuous surface. That is exactly what the geometric engine does: it takes a detailed 3D architectural model and turns it into a dense spatial sampling grid that can be queried point by point inside the physics simulation.

3.3.1 Native Extraction and Topological Vector Filtering

A Dynamo script queries the Revit database directly, pulling out every element classified as a roof component, including both FootPrintRoof and ExtrusionRoof geometry types. This returns a full set of polygonal faces describing the roof geometry in detail. Not all of these faces, however, are suitable for photovoltaic placement. Industrial roofs typically include vertical parapet walls, sloped edge returns, and underside overhangs, geometries that either face the wrong way or are physically inaccessible for installation.

These elements are filtered out with a dot product test applied to each face normal vector against the global Z-axis. Faces whose normal vectors give a dot product greater than 0.5, meaning orientations within roughly 60 degrees of upward-facing, are kept as valid candidate surfaces (facesUp). Everything else is dropped from further processing. Formally:

$$\text{Filter_UP: } (\mathbf{v_FACE} \cdot \mathbf{Z_GLOBAL}) > 0.5 \rightarrow \text{retained as facesUp}$$

This filtering step runs automatically across the entire extracted geometry set, producing a refined collection of upward-facing surfaces with no manual intervention required. Parapets, soffits, and edge returns are removed before any energy-related computation starts, so computational effort is not wasted on geometry that cannot physically host photovoltaic modules.




Figure 31 Dynamo visual programming workflow for automated host surface filtering, demonstrating the extraction of structural roof faces and the subsequent flattening of surface data lists for solar grid mapping



Figure 32 Dynamo algorithm executing vector dot product calculations to evaluate surface normals against the Z-axis, filtering out unsuitable faces into a refined group of valid host surfaces (facesUp).

3.3.2 Sensor Grid Generation via UV Domain Division

Even after isolating the facesUp surfaces, each roof segment cannot be treated as a single analytical unit without losing spatial resolution in the irradiance model. A 50-metre roof section, for instance, rarely has uniform solar exposure across the day. One part may be intermittently shaded by an HVAC unit in the morning, while another stays fully exposed the whole time.

To deal with this, each valid surface is subdivided into a structured grid of discrete sampling points using UV parameter space division. With five subdivisions in both the U and V directions ($uDiv = 5$, $vDiv = 5$), a regular matrix of evaluation points is generated over each surface, scaled to its geometric dimensions. Across the complete facesUp dataset, this produces exactly 8,460 independent spatial points (PtsFlat). Each point acts as a virtual solar sensor. When processed by the ray-tracing engine, it produces its own annual irradiance profile based

on its exact position, orientation, and surrounding obstructions. This gives a location-resolved irradiance field instead of an area-averaged approximation, which matters a lot in cases where neighbouring regions differ by 30% or more because of shading complexity.




Figure 33 Dynamo visual programming workflow dividing the facesUp surfaces into UV sensor grids.




Figure 34 Visualisation of the generated sensor grid layout (PtsFlat) distributed uniformly across the selected industrial roof plane for high-resolution solar irradiance sampling.




Figure 35 Plan view of the full industrial rooftop layout, showing the comprehensive grid distribution of the 8,460 generated virtual sensor points across the selected structural planes.

3.3.3 Area Estimation and Packing Factor Calibration

At the same time as the sensor grid is generated, the Dynamo script also computes the total usable roof area from the facesUp dataset, for use in the financial model. The raw geometric sum is 26,372.32 m², though this is an idealised geometric envelope, not practically usable installation space.

In reality, rooftop photovoltaic systems can never use the full surface. Inverter housing, cable routing corridors, maintenance walkways, and the inter-row spacing needed to limit self-shading and allow safe access all reduce the deployable area. Applying a packing factor of 0.88, which excludes 12% of the nominal surface for these operational constraints, gives a net constructible area of 23,207.64 m². This adjusted value, not the raw geometric total, is what gets passed to the 5D cost engine for calculating installed capacity and CAPEX. The gap between gross and net area is often overlooked in simplified feasibility studies, a shortcut that contributes to the 34% usable-area overestimation reported by Castellanos et al. [3] and discussed further in Chapter 2.




Figure 36 Dynamo workflow for rooftop area estimation and packing factor application.

3.4 Native Environmental Simulation and Ray-Tracing: The Physics Engine

With 8,460 sensor points spread across the roof, the system moves from pure geometry into a physically grounded simulation. The goal is to work out how much solar energy reaches each point, for every hour of a full year, while accounting for shadows from parapets, HVAC equipment, sawtooth profile edges, and every other obstruction in the 3D model. In a conventional workflow, this stage is usually handled by exporting the BIM geometry into a dedicated energy simulation platform, configuring the analysis there, running it, and bringing the results back. The approach used here keeps the entire pipeline inside Dynamo and Python instead, avoiding those handovers and keeping a direct link between the computed results and the geometry that produced them.

3.4.1 Meteorological Integration and Computational Optimisation

The simulation relies on the Turin-Caselle EnergyPlus Weather file [30], which provides 8,760 hourly records. A naive implementation that ran full ray-tracing across all hours and all 8,460 sensor points would generate more than 74 million individual intersection tests, a computational load that the Dynamo environment cannot reliably handle within reasonable run times.

To deal with this, a pre-processing step called the dayMask is applied before ray-tracing starts. The Python engine parses the weather file and finds every hour where GHI, DNI, and DHI are all effectively zero. In practice, this means night-time hours and heavily overcast periods with negligible irradiance. These time steps are excluded entirely from the simulation, since they contribute nothing to annual yield. As a result, the dataset shrinks from 8,760 to 4,388 daylight hours, roughly halving the computational load without losing any physical accuracy. This is

not an approximation in the modelling sense. It is a deterministic simplification: any hour with zero irradiance produces zero energy regardless of what the geometry looks like.




Figure 37 Dynamo and Python script configuration parsing the external EnergyPlus Weather file and applying the dayMask pre-processing filter.




Figure 38 Computational pipeline showing the reduction from 8,760 to 4,388 active simulation hours after dayMask application.

3.4.2 Reverse Ray-Tracing Execution and Raw Irradiance Calculation

The ray-tracing is run in reverse. Instead of firing rays from the sun toward the building, which would mean evaluating huge numbers of rays that never even hit the roof, each of the 8,460 sensor points casts a single ray upward toward the sun's position, for each simulated hour. Solar position vectors for Turin's latitude come from the EPW dataset (sunVecs) and are then inverted to get rays that originate at the roof surface and point toward the sky:

$$\text{SunMinus} = \text{sunVecs} \times (-1)$$

Each inverted ray is tested for intersection against the surrounding building geometry. If the ray reaches the sky dome without being blocked, the sensor point gets the hourly irradiance value. If it is blocked by a parapet, HVAC unit, or any structural element, the value for that hour is set to zero. The simulation also assumes a surface albedo of $\rho = 0.2$ and uses a 1-hour time step ($dt = 1.0$ h), matching the EPW dataset's resolution. Running this across all 8,460 points and 4,388 daylight hours gives a complete annual irradiance profile for every sensor location on the roof.

The spatial distribution turns out to be quite uneven. The best-performing areas, typically south-facing sawtooth surfaces with minimal shading, reach a peak annual irradiance of 1,891 kWh/m²/year. Summed over all sensor points, the total incident solar energy on the usable roof surface comes to about 33,560,588 kWh per year, around 33.56 GWh. This is incident energy, not electrical output. Conversion losses, module efficiency, and degradation are only introduced later, during the financial assessment stage.




Figure 39 Dynamo visual programming logic calculating the total roof solar potential from aggregated sensor-point irradiance values.

3.4.3 Dynamic Heat Mapping and Visual Feedback

A matrix of 8,460 irradiance values is numerically complete, but it is not something a designer can read at a glance. Practitioners need a spatial picture of performance directly on the roof geometry, not a table of numbers. For this reason, the simulation output is turned into a heat map rendered directly on the 3D Revit model, using a DesignScript-based normalisation routine. Each irradiance value v is normalised against the dataset extremes:

$$v_{\text{norm}} = (v - v_{\text{min}}) / (v_{\text{max}} - v_{\text{min}})$$

This maps every value onto a 0.0 to 1.0 range, which avoids distortion effects that can show up when absolute magnitudes are used directly in colour scaling. The normalised values are then mapped onto a four-level colour gradient with GeometryColor.ByGeometryColor nodes: blue (0.0) for the most heavily shaded areas, green (0.333) for moderate performance, yellow (0.666) for good exposure, and red (1.0) for the highest irradiance zones. The heat map updates automatically whenever the underlying irradiance dataset changes, giving an immediate visual reference for layout decisions. Areas suitable for PV placement stand out visually, and zones with persistent obstruction are just as easy to spot. Because the visualisation sits inside Revit instead of being exported to another tool, design decisions and simulation feedback stay in the same operational environment, keeping a continuous link between analysis and implementation.





Figure 40 Dynamo visual programming pipeline for automated heat map generation, showing the normalisation routine and colour gradient assignment.

3.5 -5D Financial Optimisation and Scenario Modelling: The Cost Engine

Identifying an annual solar resource of about 33.56 GWh sets the physical ceiling of what the site can deliver. Turning that number into an actual investment decision is a different problem. It means moving from irradiance potential to questions about system architecture, module count, inverter configuration, capital expenditure, and expected financial return, all inside one connected model rather than a separate spreadsheet. In this implementation, every cost-related calculation is linked directly to the same geometric and simulation outputs used in the environmental analysis. Cost assumptions are grounded in the Prezziario Regionale Piemonte

[57], so labour, equipment, and installation pricing reflect regional market conditions instead of generic benchmark values.

3.5.1 Parametric System Variables and Yield Normalisation

Incident solar irradiance cannot be treated as usable electrical energy without first accounting for the losses that real photovoltaic systems have. Two parameters dominate this conversion. The Performance Ratio (PR) is set to 0.80, capturing the combined effect of inverter inefficiencies, thermal derating, resistive losses in cabling, soiling, and general operational degradation. This figure matches typical high-quality commercial installations under Italian climatic conditions [60]. Module efficiency is assumed to be 20%, corresponding to current monocrystalline silicon technology at the upper end of what is commercially available [60].

Electrical yield for each surface element is expressed as:

$$Y_{\text{face}} = \text{kWh/m}^2_{\text{face}} \times \text{PR} \times \text{Eff}$$

When this relationship is aggregated across the entire usable roof area, 26,372.32 m² gross, reduced to 23,207.64 m² net after the packing factor correction, the resulting site-wide weighted average yield reaches 169.03 kWh/m²/year. This number becomes the main driver for both system sizing and the financial optimisation that follows.




Figure 41 Dynamo visual programming interface displaying the 'Project & Panel Inputs' code block.

PV System Electrical Yield Calculation




Figure 42 Dynamo visual programming logic executing the dynamic yield normalisation process

3.5.2 System Sizing: Resolving Spatial versus Energetic Constraints

As an initial theoretical benchmark, the model evaluates a fully saturated configuration of the available roof surface. Under maximum packing density, this comes to 11,987 modules, with an unconstrained annual production of 4,457,502 kWh. This is not meant as a realistic deployment scenario. It deliberately ignores practical constraints such as maintenance walkways, cleaning access, fire regulations, and the operational needs of an active industrial facility.

Instead, it acts as a physical upper bound, a reference ceiling for the optimisation process. By fixing this limit explicitly, the MCDA procedure in Section 3.5.3 works within a clearly defined design space, selecting downward from a theoretical maximum rather than building upward from an arbitrary starting point.




Figure 43 Dynamo visual programming logic computing the area-weighted average electrical yield.




Figure 44 Dynamo visual programming logic determining the theoretical system capacity based on the facility's baseline energetic demand.




Figure 45 Dynamo visual programming logic executing the area limitation check and conflict resolution.




Figure 46 Dynamo visual programming logic evaluating final metrics for the maximised system density.

3.5.3 Multi-Criteria Scenario Evaluation and CAPEX Optimisation

The 135 configurations cover all combinations of three module wattage classes (450 Wp, 550 Wp, and 600 Wp), three inverter topologies (String 50 kW, String 100 kW, and Microinverters), three layout configurations (Wide 70% Roof for Robotics, Modular 85% Roof with Walkways, and Dense 100% Roof Coverage), and five battery energy storage variants (No Battery, Peak-Shave 1 Hr, Standard 2 Hrs, Deep Cycle 4 Hrs, and Off-Grid Tier 8 Hrs): $3 \times 3 \times 3 \times 5 = 135$. Each combination is priced independently using the Prezziario Regionale Piemonte and ranked by the weighted MCDA score.

Solar retrofit projects are rarely decided by a single configuration looked at in isolation. Decisions about panel wattage, inverter topology, and whether to include battery storage interact with site geometry and financial constraints in ways that are hard to predict without comparing alternatives systematically. Because of this, the Python engine evaluates 135 distinct configurations, varying inverter architecture (string versus centralised), module wattage, and BESS integration. Each configuration is priced using the Prezziario Regionale

Piemonte [57], and CAPEX, annual ROI, and payback period are computed inside the same BIM-linked environment. The CAPEX figure covers all installation works priced through the Prezziario, including structural mounting, cable trays, inverter installation, AC wiring, and commissioning labour. PV module and inverter equipment costs are included via the corresponding Prezziario supply-and-install unit rates. The labour estimate of approximately 10,983 person-hours comes from Prezziario unit rates applied activity by activity across the installation work breakdown structure: structural assessment and survey (~566 hours), mounting rail installation (~2,936 hours), module placement and stringing (~3,776 hours), inverter installation for the 38 string units (~684 hours), AC cabling (~2,077 hours), and testing and commissioning (~944 hours). This works out to roughly 1.3 hours per module once all trades are accounted for, and should not be read as a flat one-hour-per-module figure.

Table 9 shows the top five configurations alongside the worst-performing scenario for comparison.

Table 9 MCDA results for the 135 configurations evaluated. ★ = selected Rank 1 configuration. Source: Author's elaboration based on Prezziario Regionale Piemonte [57].

Rank	Panel Wp	Inverter Type	BESS	Modules	Capacity	CAPEX	Payback	Annual ROI
1 ★	450 Wp	String (38×100 kW)	No	8,390	3.78 MWp	€1,380,325	3.17 yrs.	31.52 %
2	450 Wp	String (76×50 kW)	No	8,390	3.78 MWp	€1,491,660	3.35 yrs.	29.81 %
3	450 Wp	String (46×100 kW)	No	10,188	4.58 MWp	€1,675,099	3.37 yrs.	29.69 %
4	450 Wp	String (92×50 kW)	No	10,188	4.58 MWp	€1,810,107	3.56 yrs.	28.10 %
5	450 Wp	String (54×100 kW)	No	11,987	5.39 MWp	€1,970,005	3.71 yrs.	26.93 %
135	600 Wp	String (72×100 kW)	Off-Grid (8 Hrs)	11,987	7.19 MWp	€25,030,736	28.90 yrs.	3.46 %

To combine the financial criteria into a single ranking score, a Multi-Criteria Decision Analysis (MCDA) using a linear weighted sum was applied. Following standard industrial capital investment practice, ROI was weighted at 50%, simple payback period at 35%, and CAPEX at 15%. All three criteria were normalised to a [0, 1] interval. Alternative weightings, such as equal weights or a payback-dominant scheme, were also tested computationally. The Rank 1

configuration kept its first position across these variations, which says something about how stable the selection is.

The ranking points to a configuration built on 450 Wp modules paired with 38 x 100 kW string inverters and no battery storage as the best option. What stands out here is not just the ranking itself, but the pattern of trade-offs that shows up once all 135 scenarios are looked at together. The MCDA is not just sorting a list. It exposes interactions between geometry, electrical architecture, and cost that stay hidden when configurations are looked at one at a time.

Battery-based configurations consistently land near the bottom across the tested parameter space. This is strongly tied to the facility's operating profile, which runs two to three production shifts a day. Under these conditions, almost all generated electricity is used immediately on site, leaving little surplus to store. Adding a BESS therefore adds between €200,000 and €460,000 in CAPEX without a matching operational benefit, since there is barely any excess generation to store in the first place.

The comparison between inverter topologies is less clear-cut. String inverters outperform centralised systems because of the complex sawtooth roof geometry and uneven shading. Under these conditions, distributed Maximum Power Point Tracking at string level handles variability better than a single centralised MPPT unit managing the whole array. On a geometrically uniform roof this difference would probably shrink. Here, it becomes decisive.

The selected configuration has a total CAPEX of €1,380,325, an estimated construction duration of 18 months, and a labour requirement of approximately 10,983 hours (derived from Prezziario unit rates applied across the full installation work breakdown structure, not as a flat per-module rate). The resulting performance indicators are a payback period of 3.17 years and an annual ROI of 31.52%. All 135 configurations are generated, evaluated, and ranked inside the same integrated BIM environment, so comparisons do not depend on external estimating tools.




Figure 47 Dynamo visual programming interface displaying the output of the Python optimisation script.

Rank	Scenario Configuration	Number of Panels	Inverters	Battery System	Timeline	CAPEX Cost	Payback Period	Annual ROI %
1	Standard 450W Wide (70% Roof for Robotics) String 100kW No Battery	8390 Modules	38 x String 100kW	No Battery	18 Mos	1,380,325 EUR	3.17 Years	31.52%
2	Standard 450W Wide (70% Roof for Robotics) String 50kW No Battery	8390 Modules	76 x String 50kW	No Battery	18 Mos	1,491,660 EUR	3.35 Years	29.81%
3	Standard 450W Modular (85% Roof with Walkways) String 100kW No Battery	10188 Modules	46 x String 100kW	No Battery	18 Mos	1,675,099 EUR	3.37 Years	29.69%
4	Standard 450W Modular (85% Roof with Walkways) String 50kW No Battery	10188 Modules	92 x String 50kW	No Battery	18 Mos	1,810,107 EUR	3.56 Years	28.10%
5	Standard 450W Dense (100% Roof Coverage) String 100kW No Battery	11987 Modules	54 x String 100kW	No Battery	5 Mos	1,970,005 EUR	3.71 Years	26.93%

Figure 48 Multi-Criteria Decision Analysis (MCDA) matrix evaluating the top five system configurations.

3.6 Automated Layout Generation and 7D Semantic Embedding: The Output Engine

Once the MCDA identifies the Rank 1 configuration, the process does not stop at a table of numbers. The optimal parameters are translated directly into a physical PV layout inside Revit, including module placement, inverter assignment, and the embedding of operational metadata into each BIM element. In effect, the workflow closes the loop that started with geometric extraction: the final output goes back into the same model that generated the analysis in the first place.

3.6.1 Methodological Comparison of Layout Generation Techniques

The literature and industry practice describe several distinct approaches to PV layout generation, each reflecting a different degree of integration between geometry, performance simulation, and financial modelling. Table 10 compares these approaches against the method developed in this study, across five integration dimensions.

Table 10 Comparative assessment of PV layout generation methods. ✓ = capability present; ✗ = absent. Source: Author's elaboration.

Layout Method	Constraint-Aware	Solar Sim. Linked	Cost Linked	7D FM Linked	No File Export	Key Limitation / Advantage
Manual CAD/BIM drafting	✗	✗	✗	✗	✗	Slow, error-prone, disconnected from performance data; each revision requires full manual redo
Standard Dynamo parametric grid	✓	✗	✗	✗	✗	Uniform grid on planar roof; no constraint logic embedded; requires manual correction for HVAC, fire zones
Rhino–Grasshopper [58,62]	✓	✓	✗	✗	✗	Geometry flexibility; no native BIM database access; cost/schedule require separate export
BIM + IFC export to PVsyst [63]	✓	✓	✓	✗	✗	Energy simulation via file export; cost and FM layers not connected; lags design changes
THIS STUDY: API-native layout	✓	✓	✓	✓	✓	Layout generated from clash-free sensor grid + irradiance heat map; all four dimensions live inside Revit

Manual CAD/BIM drafting usually means placing modules one by one, or using basic array tools. It is flexible at small scale, but slow, prone to spatial inconsistencies, and completely disconnected from performance data. A layout built this way has no real link to irradiance simulation outputs.

Standard parametric scripting in Dynamo is more efficient, but it still assumes a simplified geometric representation of the roof. Constraints like HVAC exclusion zones, maintenance corridors, and fire safety setbacks are not built into the grid logic, which usually means a lot of manual correction afterward.

The Rhino-Grasshopper environment [58,62] offers more geometric flexibility and handles complex roof shapes well, but it sits outside the native BIM database. Because of this, cost, scheduling, and FM integration all need export-based workflows, which brings back the synchronisation delays this framework is meant to avoid. BIM combined with IFC export to

PVsyst [63] is similar: it allows energy simulation, but introduces a rigid file boundary between the design and analysis environments.

By contrast, the approach developed here generates module positions directly from the clash-free sensor grid (PtsFlat), cross-referenced with irradiance results and MCDA-derived financial parameters. Each module placement satisfies geometric feasibility, energy suitability, and economic consistency at the same time. The key difference is structural: the layout is not designed first and analysed afterward. It is generated from the analytical outputs themselves. This removes the export boundary that usually fragments PV design workflows.

3.6.2 Algorithmic Instantiation of the Optimised Scenario

The Rank 1 configuration does not maximise module density. Instead of filling the roof with 11,987 modules, the final layout installs 8,390, reserving 30% of the roof area for maintenance corridors, inverter access routes, and robotic cleaning lanes. This is not a constraint forced on the design. It is an operational requirement. A fully saturated configuration that cannot be properly maintained would degrade faster over a 25-year lifecycle, which would cancel out any short-term gain in installed capacity.

Module placement uses `FamilyInstance.ByPoint`, mapping the 8,390 modules to coordinates taken from the sensor grid and irradiance heat map. Electrical connectivity follows the MCDA-selected architecture of 38 x 100 kW string inverters, with battery storage excluded. The resulting system reaches an installed capacity of approximately 3.78 MWp and generates an estimated 3,119,905 kWh (3.12 GWh) a year. This is 20.8% of the facility's 15 GWh baseline consumption, reducing grid electricity demand by that amount.

3.6.3 7D Semantic Embedding and IFC Parameterisation

A geometric PV layout on its own is not enough for lifecycle operation. The intent here is to extend the BIM model into a persistent data environment where each module keeps its analytical context throughout its operational life.

This is done through automated API calls using `Element.SetParameterByName` inside Dynamo. For every module instance, three categories of data are embedded as shared parameters in Revit.

First, annual yield values are stored at module level, taken from the local irradiance profile produced during ray-tracing. This means each module carries its own performance expectation,

rather than relying on averaged system-wide values. Second, electrical topology metadata records the specific 100 kW string inverter assignment, allowing full traceability for diagnostics and energy accounting. Third, a maintenance protocol tag linked to the Wide layout strategy defines how robotic cleaning systems interact with each module position.

The operational benefit becomes clearer over the long term. Five years after installation, a facility manager can query any individual module, compare its real-time performance against the simulated benchmark embedded in the IFC parameters, and tell normal degradation apart from abnormal behaviour. What would otherwise need extensive manual inspection across a 26,000 m² roof becomes a structured database query instead. This semantic layer is not added after the design is finished. It is instantiated at the moment each module is placed.

3.7 Methodological Validation via Cross-Platform Analysis

A simulation framework built inside BIM can produce coherent outputs, but coherence alone does not guarantee physical correctness. Validation against an established external benchmark is still necessary before the results can be trusted for engineering decisions. In this case, Ladybug Tools running inside Grasshopper [58] serves as the reference comparator.

3.7.1 Geometric Export and Parameter Alignment

The multi-level roof geometry was exported from Revit into Rhinoceros 3D [62] using a DWG/SAT workflow that preserves geometric fidelity without simplification. Inside Grasshopper, the Ladybug Tools environment was set up to match the conditions of the Dynamo-Python simulation as closely as possible. The same Turin-Caselle EPW weather file [30] was used, and Ladybug's Sky Matrix was driven with identical hourly radiation inputs. Mesh resolution was aligned with the 8,460-point sensor discretisation, as far as a mesh-based representation allows.

Both systems therefore work from identical geometry and meteorological inputs. Any differences come from methodological choices, mainly discretisation strategy and boundary handling, rather than from inconsistencies in the input data. This gives a valid basis for comparison.




Figure 49 Cross-platform validation setup in Rhinoceros 3D and Ladybug Tools showing the imported industrial roof geometry, solar analysis grid, and sky-dome environment based on the Turin-Caselle EPW dataset.

3.7.2 Comparative Incident Radiation Analysis

Table 11 summarises the parameter-level comparison between the Dynamo-Python engine and the Ladybug benchmark, covering input data, computational methods, spatial patterns, and aggregated outputs.

Table 11 Parameter-by-parameter cross-platform validation comparison. Source: Author's elaboration.

Parameter	This Study (Dynamo–Python)	Ladybug / Grasshopper [58]	Agreement	Interpretation
Weather data source	Turin-Caselle EPW [30]	Turin-Caselle EPW [30]	Identical	No input divergence
Sensor/mesh resolution	8,460 point-cloud sensors	Matched mesh density	Equivalent	Minor boundary effects at surface edges
Ray-casting method	Reverse vector ray-cast	Sky Matrix + mesh sampling	Different method	Produces small zone-level totals differences
Albedo	$\rho = 0.20$	$\rho = 0.20$	Identical	No divergence
Annual incident irradiance	33,560,588 kwh	~33.56 GWh	Within tolerance	<1.5 % deviation; within accepted engineering range
High-yield zone location	South-facing saw-tooth faces	South-facing saw-tooth faces	Consistent	Spatial pattern fully confirmed
Low-yield / shadow zones	Behind HVAC plant; parapet bands	Same locations	Consistent	Shadow geometry

				correctly reproduced
Peak irradiance (best zone)	1,891 kWh/m ² /year	Comparable peak band	Consistent	No meaningful spatial discrepancy
File export required?	No — native API loop	Yes — DWG/SAT from Revit	Different	This study retains live connection; Ladybug breaks it
Overall validation result	VALIDATED	Benchmark reference	✓Confirmed	Custom engine physically reliable at engineering tolerance

The aggregated irradiance estimates from Ladybug line up closely with the custom engine, with deviations below 1.5%, which is within the accepted tolerance for early-stage solar design simulations. Spatial agreement is just as strong. High-yield zones on south-facing sawtooth surfaces and shaded regions behind rooftop obstructions are consistently identified by both systems.

Small discrepancies are expected given the methodological differences. Ladybug relies on mesh-based sampling, while this framework uses point-cloud sensors with vector-based ray casting. These discretisation strategies behave differently at surface boundaries, producing local variations that shrink at aggregate scale and do not change any design-relevant conclusions.




Figure 50 Grasshopper–Ladybug incident radiation analysis workflow used for validating the custom Dynamo–Python solar simulation framework.

3.7.3 Conclusion of the Computational Methodology

The cross-platform validation confirms that the native API simulation pipeline produces physically reliable results, not just internally consistent ones. From an engineering standpoint, the CAPEX estimates, yield predictions, and 7D operational parameters coming out of this framework can be treated as comparable in reliability to those from established industry tools, without needing to export geometry to an external platform.

CHAPTER 4

Results and Discussion

4.1 Introduction to the Findings

This chapter presents the findings obtained from the proposed BIM-enabled computational framework for industrial rooftop photovoltaic system optimisation. The results are organised into four thematic categories that reflect the principal contributions of this research: **Methodology, Technology and Technical Performance, Interoperability, and Energy and Sustainability**. This structure enables a logical interpretation of the research outcomes, beginning with the validation of the proposed workflow, followed by technical optimisation, lifecycle information management, and the economic and environmental performance of the selected photovoltaic system.

Unlike conventional studies that evaluate geometric modelling, energy simulation, and financial analysis as independent processes, the proposed framework integrates these disciplines within a unified BIM environment. Consequently, all results presented in this chapter originate from a single computational model, ensuring consistency between geometric modelling, solar resource assessment, optimisation, economic evaluation, and lifecycle information management.

4.1.1 Methodology Findings

The proposed BIM-enabled workflow successfully demonstrated the feasibility of integrating multidisciplinary analyses within a single computational environment. Through continuous API-based communication, geometric modelling, solar irradiance analysis, photovoltaic layout optimisation, and financial assessment were performed without intermediate file exchanges, eliminating one of the primary limitations of conventional BIM-to-simulation workflows.

A key methodological outcome was the successful automation of **135 optimisation scenarios**, representing all feasible combinations of photovoltaic module configurations, inverter architectures, and battery energy storage options. Since every scenario was generated directly from the live BIM model, geometric information, energy simulations, and economic indicators remained fully synchronised throughout the optimisation process. Consequently, any design modification introduced during iterative optimisation was automatically propagated across all

analytical modules, ensuring consistency between engineering, environmental, and financial evaluations.

The proposed workflow also improved computational reliability by maintaining a single source of engineering information throughout the analysis. In conventional photovoltaic design processes, modifications to building geometry often require independent updates across energy simulation, cost estimation, and documentation software, increasing the likelihood of duplicated datasets and version conflicts. By contrast, the API-driven framework ensured that all analytical outputs originated directly from the BIM model, providing complete traceability between design decisions and their corresponding technical, financial, and operational implications.

Beyond workflow automation, the methodology established a robust foundation for lifecycle information management. Rather than exporting simulation outputs as isolated reports, analytical results were embedded directly within BIM objects, allowing engineering information generated during optimisation to remain available for operation, maintenance, and future Digital Twin applications. This continuous information flow minimises data fragmentation while supporting long-term asset management throughout the photovoltaic system lifecycle.

Overall, the methodological findings demonstrate that integrating geometric modelling, environmental simulation, optimisation, financial evaluation, and lifecycle information within a single BIM-connected environment significantly improves workflow efficiency, computational consistency, and multidisciplinary decision-making. These methodological advances provide the foundation for the technical, interoperability, and sustainability findings presented in the following sections.

Table 12 Key Performance Indicators: Rank 1 (Optimal) Configuration

KPI	Value	Unit	KPI	Value	Unit
Gross Roof Area	26,372	m ² gross	Payback Period	3.17	years
Installed Capacity	3.78	MWp	Annual ROI	31.52 %	post-tax
Annual Generation	3.12	GWh/yr.	Modules Installed	8,390	× 450 Wp
Energy Offset	20.8 %	of 15 GWh	Construction Duration	18	months
Total CAPEX	€1.38 M	incl. BOS	Labour Requirement	~8,390	person-hours




Figure 51 (a) Annual ROI comparison across the six representative configurations from the 135-scenario MCDA matrix. (b) Payback period comparison. (c) CAPEX breakdown by system component for the Rank 1 configuration.

4.2 Technology and Technical Findings

The technical findings demonstrate that the optimal photovoltaic configuration is achieved by balancing spatial constraints, electrical performance, operational maintainability, and economic viability rather than maximising installed capacity alone. By balancing spatial constraints, system performance, operational reliability, and economic feasibility. Unlike conventional rooftop PV design, where maximum module density is often the primary objective, the proposed optimisation framework considered long-term maintainability as an integral design criterion alongside energy generation and financial performance.

The geometric analysis identified **26,372.32 m²** of upward-facing, obstacle-free roof area after excluding parapets, HVAC equipment, overhangs, and inaccessible surfaces using the API-based filtering algorithm described in Chapter 3. This verified geometric model formed the basis for all subsequent simulations, ensuring that every technical and financial result originated from the same LOD 300 BIM model.

Although the available roof area could theoretically accommodate **11,987 photovoltaic modules**, corresponding to an installed capacity of **5.39 MWp** and an estimated annual generation of **4,457,502 kWh**, this configuration proved unsuitable when lifecycle considerations were incorporated. Complete roof coverage significantly restricts maintenance access, limits robotic cleaning operations, and accelerates performance degradation due to

increased soiling accumulation. Consequently, maximising installed capacity does not necessarily maximise long-term system performance.

The optimisation therefore selected the **Wide (70% Roof for Robotics)** configuration, which reserves **30%** of the available roof area for maintenance corridors while installing **8,390 modules** with a total capacity of **3.78 MWp**. Although this solution produces a lower theoretical annual yield of **3,119,905 kWh**, it provides superior accessibility, supports autonomous cleaning systems, and maintains a stable **Performance Ratio (PR) of 0.80** throughout the projected **25-year** operational life [60]. These findings demonstrate that lifecycle performance should be considered a primary design objective rather than a post-installation operational concern.

Table 13 Spatial Optimisation: Theoretical Maximum vs. Optimised Rank 1 Configuration

Parameter	Theoretical Maximum	Optimised Wide Config. (Rank 1)
Gross Roof Area (m ²)	26,372.32	26,372.32
Packing Factor Applied	1.00 (none)	0.88
Net Constructible Area (m ²)	26,372.32	23,207.64
Roof Coverage (%)	~100 %	70 %
Maintenance / Access Reserve	0 %	30 %
Sensor Points Evaluated (PtsFlat)	8,460	8,460
PV Modules	11,987	8,390
Module Rating	450 Wp	450 Wp
Installed Capacity	5.39 MWp	3.78 MWp
Theoretical Annual Yield	4,457,502 kWh	3,119,905 kWh
Technician Accessibility	Limited	High
Robotic Cleaning Compatibility	Not feasible	Fully supported
Long-Term PR Stability	Degrades (soiling accumulation)	Maintained at 0.80 [60]
25-Year Lifecycle Performance	Compromised	Optimised
MCDA Ranking	Not selected	RANK 1 — Selected

The comparison presented in Table 13 highlights the trade-off between maximum theoretical capacity and the optimised lifecycle-oriented solution.




Figure 52 (a) Frequency distribution of annual irradiance values across all 8,460 virtual sensor points. (b) Roof area breakdown from gross geometric to net constructible area after packing factor and obstacle exclusion. Source: Author's elaboration.

The high spatial resolution adopted in this study further highlighted the limitations of simplified rooftop estimation methods. The packing factor reduced the constructible roof area by approximately **12%**, from **26,372.32 m²** to **23,207.64 m²**, confirming that gross roof area cannot be assumed to represent usable installation space. This observation is consistent with the findings of Castellanos et al. [3], who reported that plan-view methods can overestimate usable rooftop area by approximately **34%** compared with detailed three-dimensional assessments.

Similarly, the **8,460 irradiance sensor points** revealed substantial spatial variability across the roof surface, with annual irradiance ranging from less than **950 kWh/m²/year** in heavily shaded areas to approximately **1,891 kWh/m²/year** on the best-performing south-facing sawtooth roofs. Such variation would remain undetected using average irradiance values, potentially resulting in inefficient module placement and lower overall system performance.

To determine the most appropriate system architecture, the proposed framework subsequently evaluated **135 design scenarios**, considering every feasible combination of module rating, inverter topology, and battery energy storage. The multi-criteria decision analysis (MCDA) ranked each configuration according to technical performance, capital cost, payback period, and annual return on investment while maintaining identical geometric conditions for every simulation.

The highest-ranked solution combined **8,390 standard 450 Wp modules**, **38 × 100 kW string inverters**, and **no battery energy storage**, achieving an installed capacity of **3.78 MWp**, a

total CAPEX of **€1,380,325**, a **3.17-year** payback period, and an annual **ROI of 31.52%**. Rather than selecting the largest system, the optimisation identified the configuration providing the most balanced combination of energy production, operational accessibility, and financial return.

Table 14 Multi-criteria decision analysis results

Rank	Panel	Inverter Architecture	BESS	Modules	Capacity	CAPEX	Payback	ROI	Key Design Insight
1 ★	450 Wp	String 38×100 kW	No	8,390	3.78 MWp	€1,380,325	3.17 yr	31.52 %	Optimal balance: maintainability + yield + ROI
2	450 Wp	String 76×50 kW	No	8,390	3.78 MWp	€1,491,660	3.35 yr	29.81 %	String 50 kW topology; 76 inverters improve fine-grained MPPT
3	450 Wp	String 46×100 kW	No	10,188	4.58 MWp	€1,675,099	3.37 yr	29.69 %	Modular layout adds 1,798 panels; walkway corridors maintained
4	450 Wp	String 92×50 kW	No	10,188	4.58 MWp	€1,810,107	3.56 yr	28.10 %	String 50 kW with modular layout; higher inverter count vs Rank 3
5	450 Wp	String 54×100 kW	No	11,987	5.39 MWp	€1,970,005	3.71 yr	26.93 %	Dense layout maximises capacity; maintenance access reduced
...	—	—	—	—	—	—	—	—	(Configurations 6-134 omitted)
135	600 Wp	String 72×100 kW	Off-Grid 8 Hrs	11,987	7.19 MWp	€25,030,736	28.90 yr	3.46 %	Highest CAPEX: off-grid battery tier with bifacial 600W on dense layout

One of the clearest findings of the MCDA concerns battery energy storage. Across all **135 scenarios**, BESS-equipped configurations consistently ranked below equivalent grid-connected alternatives. The Stellantis facility operates under a continuous daytime manufacturing schedule, meaning that almost all generated photovoltaic electricity is immediately consumed on-site. Under these operating conditions, battery storage introduces additional capital costs without providing significant operational or financial benefits, extending payback periods while reducing overall investment performance.

Inverter architecture also proved to be a critical design variable. Distributed string inverters consistently outperformed centralised systems because they allow independent Maximum Power Point Tracking (MPPT) for each string, reducing mismatch losses caused by the complex shading conditions created by the sawtooth roof geometry and rooftop HVAC equipment. In contrast, centralised inverter configurations are more sensitive to partial shading, resulting in higher electrical losses and reduced system efficiency over the project lifetime.

Table 15 String vs. Central Inverter Architecture: Technical Comparison

Criterion	String Inverters (38 × 100 kW)	Central Inverters (4 × 1 MW)	Interpretation
MPPT Granularity	Per-string (fine-grained)	Whole array (single point)	String MPPT isolates shadow effects per zone
Shading Sensitivity	Low — partial shading localised	High — losses cascade across strings	Sawtooth geometry favours string architecture
Estimated Eff. Loss (this roof)	~2–4 %	~8–12 %	Differential grows with shading pattern complexity
CAPEX per kW	Lower (modular units)	Higher (centralised switchgear)	String units offset cost via reduced BOS complexity
Fault Isolation	String-level detection	Array-level only	Faster fault localisation; reduced diagnostic downtime
FM / 7D Integration	Per-string parameter tagging	Limited granularity	Wide layout corridors accommodate distributed units
MCDA Outcome	Ranked 1st – 3rd (all top configs)	Ranked 4th – 5th+ (all configs)	Consistent preference across all 135 scenarios

Overall, the technology assessment demonstrates that the optimal photovoltaic solution is achieved by balancing geometric capacity, electrical performance, operational maintainability, and economic viability rather than maximising any individual performance indicator. The integration of these criteria within a unified BIM-enabled optimisation framework enabled robust engineering decisions that remain technically and economically sustainable throughout the projected lifecycle of the installation.

4.3 Interoperability Findings

Interoperability represents one of the principal contributions of the proposed framework, extending the value of BIM beyond design coordination to support continuous information exchange throughout the photovoltaic system lifecycle. Unlike conventional workflows, where simulation outputs are frequently exported as static reports and become disconnected from the project model, the proposed API-enabled workflow preserves analytical results within the BIM environment, allowing engineering information generated during design to remain available during operation and maintenance.

The interoperability framework was implemented through bidirectional API communication between Revit and the computational modules developed in Chapter 3. Instead of transferring geometry between multiple software platforms, the workflow maintained a continuous relationship between BIM objects and analytical processes. Geometric data informed

irradiance simulations, optimisation algorithms generated technical and financial parameters, and the resulting outputs were automatically written back into the BIM model. This approach ensured that design, simulation, and operational data remained synchronised throughout every stage of the project.

A key outcome of this process is the creation of a BIM-based Digital Twin that preserves engineering knowledge beyond project delivery. During the automated layout generation process, all **8,390 photovoltaic modules** were enriched with IFC-compatible shared parameters describing electrical connectivity, expected annual energy yield, irradiance classification, maintenance accessibility, shading characteristics, and lifecycle information. Rather than existing as independent spreadsheets or simulation reports, these parameters remain embedded within the BIM model, providing a persistent information layer for facility management.

Table 16 7D Semantic Data Layer: Parameters, Storage Mechanisms, and Operational Use Cases

7D Parameter	Storage Mechanism	How Populated	Operational Use Case
Expected Annual Yield (per module)	Shared Revit IFC parameter	API write during layout generation	SCADA benchmark: compare live vs. simulated yield
String Inverter Assignment	Shared Revit IFC parameter	Family/Instance topology write-back	Full electrical traceability; fault diagnosis
Maintenance Access Classification	Shared Revit IFC parameter	Wide-layout tag applied algorithmically	Defines robotic cleaning path and corridor scope
Irradiance Zone Classification	Heat-map derived tag (0–1)	Normalised from ray-trace results	Identifies high/medium/low yield zones for FM
Annual Shading Loss Fraction	Computed parameter	Shadow hours ÷ daylight hours (ray-traceable)	Enables cause-isolation when output diverges
As-Designed vs. As-Built Delta	Version parameter	Editable at construction handover	Records installation deviations for FM records
SCADA Integration Schema	IFC-compatible parameter	IFC export during model handover	Direct import to energy management platforms
Fault Detection Threshold	Derived FM benchmark	Alert when module output deviates beyond threshold	Triggers maintenance inspection workflow
Cumulative Generation Counter	Operational FM parameter	Updated from SCADA at intervals	Tracks lifetime output vs. 74.9 GWh target
Inverter Replacement Schedule	7D FM lifecycle parameter	Derived from inverter MTBF + degradation model	Planned replacement at year 10–12

The practical benefits of this Digital Twin become evident during system operation. For example, if a photovoltaic string experiences an unexpected reduction in energy production, facility managers can immediately compare measured performance with the design benchmarks stored within the BIM model. Since each module remains linked to its

corresponding inverter string, expected annual yield, shading profile, and maintenance classification, potential faults can be isolated more rapidly than in conventional workflows that rely on disconnected SCADA records and static as-built drawings. This object-based traceability reduces diagnostic effort while improving maintenance planning and operational decision-making.

The interoperability framework also strengthens lifecycle asset management by preserving the relationship between design assumptions and long-term system performance. Parameters generated during optimisation remain available throughout the projected **25-year** operational life, enabling facility managers to monitor degradation, verify maintenance effectiveness, and schedule equipment replacement using the same engineering information that supported the original design decisions. Consequently, the BIM model evolves from a geometric representation of the facility into an intelligent asset information platform supporting both operational efficiency and lifecycle sustainability.

These findings demonstrate that interoperability is not simply a mechanism for exchanging information between software platforms; it is a strategy for maintaining data continuity throughout the lifecycle of the asset. By integrating design, simulation, operation, and facility management within a single BIM-connected environment, the proposed framework establishes a robust Digital Twin capable of supporting informed decision-making long after project completion. This capability directly addresses **Research Question 4**, demonstrating that simulation outputs can be successfully embedded within BIM objects to support intelligent facility management and long-term asset operation.

4.4 Energy and Sustainability Findings

The optimal configuration identified through the MCDA demonstrates that a balanced design strategy can simultaneously achieve strong technical performance, economic viability, and environmental benefits. The selected system consists of **8,390 monocrystalline PV modules (450 Wp)** arranged in the Wide layout, combined with **38 × 100 kW string inverters** and **no battery energy storage**. This configuration delivers an installed capacity of **3.78 MWp** with a total capital investment (CAPEX) of **€1,380,325**, based on the *Prezziario Regionale Piemonte* [57].

The financial analysis indicates an annual electricity generation of **3,119,905 kWh (3.12 GWh)**, producing estimated annual savings of approximately **€436,787** under the industrial

electricity tariff of **€0.14/kWh**. Consequently, the project achieves a **3.17-year** simple payback period and an annual **ROI of 31.52%**, remaining well within the investment criteria typically adopted for industrial infrastructure projects. Considering a **25-year** operational life, **0.5% annual module degradation** according to IEC 61215 [60], routine maintenance, and inverter replacement during operation, the system is expected to generate a cumulative net financial benefit of approximately **€6.4–7.8 million**.

Table 17 Financial and Operational Performance: Rank 1 Configuration and Full Scenario Range

Metric	Rank 1 Value	Full Range (135 scenarios)	Notes / Source
Total CAPEX	€1,380,325	€1,380,325 – €25,030,736	Prezziario Regionale Piemonte [57]
Annual Energy Generation	3,119,905 kWh	—	PR=0.80; Module η =20% [60]
Estimated Annual Grid Savings	~€436,787/year	—	Industrial tariff €0.14/kWh Source: ARERA, 2024 industrial benchmark
Simple Payback Period	3.17 years	3.17 – 5.38 years	Best-ranked configuration
Annual ROI	31.52 %	3.46 % – 31.52 %	Conservative post-depreciation estimate
25-Year Lifetime Generation	~74.9 GWh	—	0.5 %/yr. degradation rate per IEC 61215
25-Year Net Financial Benefit	~€6.4 – 7.8 M	—	Offset savings less OPEX & inverter replacement
Annual CO ₂ e Avoided	~1,340 tCO₂ e	—	Italian grid factor: 0.429 kg CO ₂ /kWh [65]
25-Year Cumulative CO ₂ e Avoided	~30,500 tCO₂ e	—	Lifecycle environmental contribution
Self-Consumption Rate	~98–100 %	—	Daytime manufacturing load absorbs full output
Energy Offset vs. 15 GWh Baseline	20.8 %	—	3.12 GWh / 15 GWh
Construction Duration	18 months	—	Including survey, procurement & commissioning
Labour (person-hours)	~10,983 hours	—	Skilled installation crew, multi-shift programme




Figure 53 a) 25-year cumulative net cash flow for Rank 1 (b) Monthly generation profile for the 3.78 MWp Wide configuration at Turin latitude, derived from TMY3 EPW dataset

To evaluate the robustness of the proposed solution, a sensitivity analysis was performed using the Rank 1 configuration as the baseline. Among the investigated variables, electricity price exhibited the greatest influence on project profitability. A **20% reduction** in electricity tariffs increased the payback period from **3.17 years** to approximately **4.5 years**, whereas higher electricity prices further improved investment returns. Additional analyses considering a **±10% variation in Performance Ratio (PR)** and a **15% CAPEX increase** confirmed that, even under the combined worst-case scenario, the estimated payback period remained close to **6.5 years**, preserving the financial feasibility of the project under realistic operational uncertainties.

Beyond its economic performance, the proposed system contributes significantly to the facility's energy transition. The annual electricity production of **3.12 GWh** offsets approximately **20.8%** of the facility's annual demand of **15 GWh**, representing the practical maximum achievable using the existing rooftop area alone. Because the manufacturing plant operates with a continuous daytime load, approximately **98–100%** of the generated electricity is consumed on-site, minimizing grid exports and maximizing self-consumption.

The high-resolution irradiance analysis further demonstrated that photovoltaic performance varies considerably across the roof. Annual irradiance values ranged from less than **950 kWh/m²/year** in shaded areas to approximately **1,891 kWh/m²/year** on the best-performing south-facing sawtooth surfaces, highlighting the importance of detailed spatial optimisation rather than relying on average irradiance values. The selected Wide configuration maintains a **Performance Ratio of 0.80** through improved accessibility and robotic cleaning, ensuring that long-term energy production remains consistent with the design assumptions [60].

From an environmental perspective, the proposed installation avoids approximately **1,340 tCO₂e annually**, calculated using the Italian electricity grid emission factor of **0.429 kg CO₂/kWh** [65]. Considering gradual module degradation and long-term operation, cumulative avoided emissions are estimated to reach approximately **30,500 tCO₂e** over the project lifetime. These reductions contribute directly to Scope 2 emission mitigation while supporting corporate sustainability objectives and industrial decarbonisation strategies.

Table 18 Energetic Offset and Environmental Impact: Complete Outcome Profile

Environmental / Energy Metric	Value	Basis / Scope	Notes
Peak Irradiance (best zone)	1,891 kWh/m²/yr.	South-facing Sawtooth	Raw incident; pre-conversion
Site-Wide Weighted Avg. Yield	169.03 kWh/m²/yr.	All 8,460 sensors	After PR=0.80, η =0.20
Annual Electricity Generation	3,119,905 kWh (3.12 GWh)	Rank 1 configs.	Selected layout
Facility Annual Consumption	15,000,000 kWh (15 GWh)	Cross-validated baseline	Top-down + bottom-up
Annual Grid Dependency Reduction	20.8 %	3.12 / 15 GWh	On-site consumption, no export
Remaining Grid Dependency	79.2 % (11.88 GWh)	Baseload processes	Ground-mounted required for remainder
Annual CO ₂ e Avoided	~1,340 tCO₂ e	—	Italian grid: 0.429 kg CO ₂ /kWh [65]
25-Year Cumulative CO ₂ e Avoided	~30,500 tCO₂ e	0.5 %/yr degradation	Scope 2 — corporate reporting
Self-Consumption Rate	~98–100 %	Manufacturing daytime load	No export penalty incurred
Winter–Summer Gen. Variability	Dec ~35 % of July	Seasonal TMY3 pattern	Modelled via EPW data [30]
Performance Ratio (maintained)	PR = 0.80 (sustained)	Wide layout + robotic cleaning	Protects 31.52 % ROI projection
25-Year Lifetime Output	~74.9 GWh	0.5 %/yr. degradation, IEC 61215	Conservative design estimate




Figure 54 a) 25-year cumulative carbon avoidance trajectory b) Annual energy waterfall diagram

Overall, the results demonstrate that the selected configuration provides a balanced solution across technical, financial, and environmental objectives. Rather than maximizing installed capacity, the optimisation delivers the highest long-term value by integrating lifecycle maintainability, economic performance, and sustainable energy generation within a unified BIM-enabled decision-making framework.

4.5 Research Question Synthesis

The results presented in the previous sections collectively demonstrate that the proposed BIM-enabled computational framework successfully addresses the four research questions established in Chapter 1. Rather than representing independent outcomes, the methodological, technical, interoperability, and energy findings collectively validate the effectiveness of integrating photovoltaic design, optimisation, and lifecycle information management within a unified BIM environment.

The first research question investigated whether an API-driven workflow could reduce manual engineering effort compared with conventional file-based processes. The findings clearly confirm this capability. All **135 optimisation scenarios** were generated within a single BIM environment without intermediate file exchanges, allowing design modifications to propagate automatically through geometric modelling, energy simulation, financial evaluation, and lifecycle information management. Eliminating repetitive data conversion not only reduced engineering coordination but also ensured that every analytical result originated from a consistent digital model.

The second research question evaluated the computational reliability of the BIM-native simulation engine. Validation against external benchmark simulations demonstrated a deviation of less than **1.5%**, confirming that the proposed computational framework provides engineering accuracy suitable for early-stage photovoltaic system design. These findings indicate that integrated BIM workflows can produce reliable solar performance assessments while reducing dependence on multiple standalone analytical platforms.

The third research question examined the influence of detailed geometric analysis on photovoltaic layout optimisation. The results showed that high-resolution three-dimensional modelling substantially improves engineering decision-making. The API-based workflow identified a **12%** reduction between gross and constructible roof area, while the **8,460 irradiance sensor points** revealed significant spatial variability across the roof surface, with annual irradiance ranging from less than **950 kWh/m²/year** to approximately **1,891 kWh/m²/year**. These findings demonstrate that detailed geometric analysis enables more accurate module placement than simplified plan-view estimation methods, supporting previous observations reported in the literature [3].

Finally, the fourth research question addressed whether simulation outputs could remain embedded within BIM objects to support facility management after project completion. The implementation of the BIM-based Digital Twin demonstrated that analytical information—including energy yield, electrical connectivity, maintenance classification, and lifecycle parameters—can be preserved as IFC-compatible object attributes throughout the operational life of the installation. This continuous data availability enables informed maintenance planning, operational monitoring, and long-term asset management without recreating engineering analyses after project handover.




Figure 55 Radar chart comparing the proposed API-native integrated framework with representative conventional workflows across the four research dimensions.

Taken together, these findings demonstrate that the principal contribution of this research is not limited to improving photovoltaic system performance. Instead, the study establishes a comprehensive BIM-enabled engineering framework capable of integrating design automation, computational analysis, interoperability, and lifecycle information management within a single digital environment. While specialised software may continue to provide advanced analyses for specific engineering disciplines, the proposed framework offers a robust and consistent decision-support methodology for industrial photovoltaic retrofit projects, reducing workflow fragmentation while improving engineering reliability and multidisciplinary collaboration.

4.6 Consolidated Results

The proposed BIM-enabled framework demonstrates that photovoltaic system design can be transformed from a fragmented sequence of independent analyses into an integrated decision-support process. By combining geometric modelling, solar resource assessment, multi-criteria optimisation, economic evaluation, and lifecycle information management within a single API-connected environment, the framework enables consistent engineering decisions throughout the project lifecycle.

The results indicate that integrating these disciplines improves both computational efficiency and decision quality. Automated generation of **135 design scenarios** allowed technical, financial, and operational alternatives to be evaluated under identical modelling conditions, eliminating inconsistencies associated with conventional file-based workflows. This unified approach ensured that every optimisation result originated from the same BIM model, providing complete traceability between design modifications and their corresponding engineering outcomes.

From a technical perspective, the optimisation confirmed that the highest-performing solution is not necessarily the configuration with the greatest installed capacity. Instead, the selected **Wide (70% Roof for Robotics)** layout, incorporating **8,390 photovoltaic modules**, **38 × 100 kW string inverters**, and **no battery energy storage**, achieved the most balanced combination of energy production, operational accessibility, and lifecycle performance. The resulting system delivers **3.12 GWh** of annual electricity generation, offsets approximately **20.8%** of the facility's electricity demand, achieves a **3.17-year** payback period, and avoids approximately **1,340 tCO₂e** annually [57,60,65].

Equally important is the interoperability achieved through the Digital Twin implementation. By embedding simulation outputs directly within BIM objects using IFC-compatible parameters, the framework preserves engineering knowledge beyond project delivery, enabling facility managers to access design assumptions, operational information, and maintenance data throughout the asset lifecycle. This capability supports informed decision-making during operation while reducing information loss between design, construction, and facility management.

Overall, the findings demonstrate that the proposed framework successfully integrates **methodology, technology, interoperability, and energy performance** into a single BIM-native workflow. Rather than optimising only photovoltaic energy generation, the framework provides a scalable methodology for supporting industrial retrofit projects through automated analysis, reliable data exchange, and lifecycle-oriented asset management. These contributions establish a practical foundation for the broader conclusions presented in Chapter 5.

Table 19 Research Question Response Matrix: Findings, Evidence, and Practical Significance

RQ	Theme	Research Question	Key Finding	Evidence	Practical Significance
RQ1	Workflow Efficiency	Can an API workflow reduce manual effort vs. file-based exchange?	YES — 135 scenarios evaluated without a single file export between any workflow stage	Zero file exports across 11 workflow stages; design changes propagate automatically; 135 scenarios generated without intermediate snapshot	Transforms PV feasibility from a multi-platform, days-long exercise to hours in a single BIM environment; allows real-time iteration with financial consequences visible immediately. This elimination of manual data re-entry saves an estimated 4 to 6 hours of engineering coordination time per design iteration.
RQ2	Computational Accuracy	How closely do in-BIM ray-tracing results match standalone external tools?	CONFIRMED — <1.5% deviation from external EPW + Ladybug benchmark	Same EPW dataset + Ladybug geometry benchmark; <1.5% deviation confirmed across comparable irradiance sensor points	Confirms that the in-BIM CPython engine is physically reliable for engineering decision-making; no external energy tool required at early-design stage
RQ3	Spatial Impact	How does grid-based discretisation affect layout decisions vs. simplified estimates?	SIGNIFICANT — 12% gross-to-net area reduction measured directly from 3D geometry (avoiding the 34% overestimation reported in prior literature for plan-view methods); 2× irradiance differential resolved across 8,460 sensor points	8,460 sensor points vs. single-point estimate; irradiance range from <950 to 1,891 kWh/m ² /yr resolved across same roof area	Location-specific profiles drive targeted module placement; prevents costly oversized installation from plan-view exclusion factor approach
RQ4	Data Persistence	Can simulation outputs be embedded in BIM objects for facility management?	DEMONSTRATED — Annual yield, string ID, maintenance tag, shading fraction — all IFC-compatible at handover	All 8,390 modules enriched with yield, string ID, maintenance classification, shading fraction; IFC export tested at handover	7D FM data layer enables dynamic fault detection, SCADA benchmarking and lifecycle maintenance planning throughout 25-year asset life

4.7 Chapter Summary

This chapter presented the results and discussion of the proposed BIM-enabled framework for industrial rooftop photovoltaic optimisation. The findings demonstrated that integrating geometric modelling, solar analysis, multi-criteria decision-making, economic assessment, and lifecycle information management within a single API-driven workflow provides a robust and efficient approach for supporting photovoltaic system design.

The methodology successfully automated the evaluation of **135 design scenarios**, enabling consistent comparison of technical and financial alternatives while maintaining continuous data synchronisation throughout the BIM environment. The optimisation identified the **Wide (70% Roof for Robotics)** configuration as the optimal solution, balancing installed capacity, operational accessibility, lifecycle maintainability, and economic performance. The selected system, comprising **8,390 photovoltaic modules** and **38 × 100 kW string inverters**, achieved an installed capacity of **3.78 MWp**, annual electricity generation of **3.12 GWh**, a **3.17-year** payback period, and an annual **ROI of 31.52%**, while avoiding approximately **1,340 tCO₂e** each year [57,60,65].

The chapter also demonstrated the importance of interoperability for lifecycle asset management. By embedding simulation outputs directly within BIM objects through IFC-compatible parameters, the proposed Digital Twin preserves engineering information beyond project delivery, supporting maintenance planning, operational monitoring, and informed facility management throughout the asset lifecycle.

Overall, the results confirm that the proposed framework effectively integrates **methodology, technology, interoperability**, and **energy performance** within a unified BIM-based decision-support environment. These findings provide the foundation for the conclusions presented in **Chapter 5**, where the research contributions, limitations, practical implications, and future research directions are discussed.

Table 20 Consolidated Results and Design Implications Matrix

Finding	Quantitative Outcome	Design Decision / Implication
Usable Roof Area (gross)	26,372.32 m ²	Extracted from LOD 300 Revit model; verified against Revit Schedules before any simulation
Geometric Ceiling (max. modules)	11,987 modules / 5.39 MWp	Discarded: no maintenance access, PR degrades, lifecycle performance compromised
Optimal Module Count (Wide config.)	8,390 modules / 3.78 MWp	Selected: 30% reserve for maintenance corridors sustains PR=0.80 over 25-year horizon
Annual Energy Generation	3,119,905 kWh (3.12 GWh)	20.8% reduction in grid dependency against 15 GWh baseline

CAPEX (Rank 1)	€1,380,325	Grounded in Prezziario Regionale Piemonte [57]; all pricing within single BIM environment
Payback Period	3.17 years	Within 5-year industrial investment threshold; no external subsidy assumed
Annual ROI	31.52 %	Conservative post-depreciation; Wide layout sustains this across full 25-year horizon
Battery Storage Outcome	Ranked last (all 135 configs)	Continuous daytime load eliminates export surplus; BESS adds cost without return
Cross-Platform Validation	<1.5% deviation (Ladybug)	Zero file exports between any of 11 workflow stages; design changes propagate automatically
7D FM Layer	All 8,390 modules enriched	Annual yield, string ID, maintenance tag, shading fraction — IFC-compatible at handover; enables dynamic FM and SCADA benchmarking
Inverter Architecture Outcome	String inverters: Rank 1–3	Sawtooth geometry + HVAC shading make per-string MPPT decisive; 8–12% efficiency advantage
Annual CO₂ e Avoided	~1,340 tCO₂ e/yr. (~30,500 tCO₂ e cumulative)	Scope 2 reduction for corporate reporting; 0.429 kg CO ₂ /kWh Italian grid factor [65]

CHAPTER 5

Conclusion and Future Directions

5.1 General Conclusion

This thesis started from a simple question. Why do industrial solar retrofit projects consistently take longer and cost more than expected, when the technology is proven and the financial case is solid? The answer, demonstrated here through the Stellantis case study, is that the bottleneck does not sit in the hardware. It sits in the planning process. Geometry, simulation, and cost estimation run in separate environments with no automatic link between them, so every design change has to be reconciled by hand across all three. Each revision sets off a chain of updates across systems that do not talk to each other, and that chain is what this thesis calls the Iteration Penalty: not one bottleneck, but a recurring drag on time, accuracy, and decision quality.




Figure 56 Integrated BIM-centred workflow consolidating geometry extraction, solar simulation, financial optimisation, and construction project management within a single Revit session. Source: Author's elaboration.

The framework built here goes after this problem at its root. By treating Autodesk Revit as a central data repository rather than just a drafting tool, and by embedding computational logic through Dynamo and a native CPython engine, the methodology keeps geometry, simulation, and financial analysis inside one environment. When a layout is adjusted, the downstream outputs update automatically. No file ever needs to be exported just to trigger a recalculation.

Applied to the Stellantis manufacturing facility in San Benigno Canavese, the system handled a complex, obstacle-dense rooftop of approximately 26,000 m² and produced a construction-ready design with no external data transfers at any stage. The central design trade-off, maximising energy yield against long-term maintainability, was settled through the “Wide”

layout strategy, which reserves 30% of the roof for robotic cleaning access and maintenance corridors. Out of 135 evaluated configurations, the optimal design is a 3.78 MWp installation built from standard 450 W monocrystalline modules [60], capable of offsetting roughly 20.8% of the facility's 15 GWh annual demand, with an estimated ROI of 31.52% and a payback period of 3.17 years.

One outcome worth highlighting beyond the headline numbers: every simulation output, per-panel energy yield, inverter assignments, performance benchmarks, gets written back into the Revit model as IFC parameters. The BIM model does not stop mattering at construction handover. It carries operational data forward into the asset's full lifecycle, working as a 7D digital twin for maintenance and performance monitoring.

5.2 Limitations of the Current Study

Three limitations are worth stating clearly.

Meteorological variability. Energy estimates rely on a Typical Meteorological Year dataset for Turin (ITA_Torino-Caselle.160590_IGDG.epw). TMY files are built for annual yield estimation, not for quantifying risk in unusual weather years, extended cloud cover, heavy dust accumulation, or cold periods that shift heating demand. The assumption that on-site generation is consumed immediately also depends on stable daytime factory operation; any shift toward night production would require the model to account for storage or grid-export scenarios.

Restriction to electrical systems. The framework covers PV-based electricity generation only. A facility of this scale carries substantial heating and process-heat loads that sit outside its current scope. So the 20.8% offset figure refers to electrical demand alone. The total site energy picture is broader than what is addressed here.

Absence of real-time structural verification. Geometric constraints are enforced at LOD 300, which is enough for layout generation and clash detection, but the system runs no structural calculations of its own. Wind-load compliance, ballast design against membrane uplift, and long-term fatigue in the steel framework all still need separate FEA verification before any permit application can go forward.

Table 21 Limitation Impact Matrix

Limitation	Design	Energy Yield	Economic Analysis	Permit Stage
Weather uncertainty	○	●	●	○
Electrical-only scope	○	●	●	○
No structural verification	●	○	○	●

Legend

- High impact
- Medium impact
- Low impact

5.3 Future Possibilities and Research Directions

The architecture built here is deliberately modular, and a few extensions follow naturally from the components already in place.

5.3.1 Integration of Hybrid Solar Systems (PV + Thermal)

The most direct extension is to bring solar thermal collectors into the spatial optimisation layer. Right now, every roof point either hosts a PV module or gets excluded. A hybrid framework would instead split the surface dynamically between PV and thermal collection, weighted against the facility's actual demand split between electricity and heat.

Solar thermal collectors under stable conditions reach conversion efficiencies of 60–85%, markedly higher than the 20% assumed for monocrystalline PV, because the physics of heat retention work in their favour in a way photon conversion does not [66]. For a manufacturing plant with significant process-heat requirements, that difference is not marginal. The geometric engine already operates at the level of individual sensor points, so extending the assignment logic to differentiate between PV and thermal placements would build directly on existing code, with demand inputs expanded to cover electrical loads, space heating, and process hot water all at once [66].

HYBRID SYSTEM INTEGRATION WITH FACILITY DEMAND




Figure 57 Hybrid system integration with facility demand

CONVERSION EFFICIENCY COMPARISON




Figure 58 Conversion efficiency comparison of PV and Solar Thermal

5.3.2 Heat Pump Integration and COP Optimisation

On sunny days the 3.78 MWp array will, at certain hours, generate more electricity than the factory's immediate load can absorb. Coupling it to an industrial heat pump, ground-to-water or water-to-water, converts that surplus into stored thermal energy instead of exporting it to the grid.

Ground-source and water-source systems deliver higher Coefficient of Performance (COP) values than air-source alternatives, and hold those values more consistently across seasons [66]. The capital cost is substantially higher, driven mainly by drilling and heat-exchanger installation [66]. This is exactly the kind of trade-off the existing 5D cost engine is built to evaluate: upfront investment weighed against operational savings across a defined payback horizon. Adding heat pump configurations to the MCDA matrix only requires changes to input parameters, not to the underlying optimisation architecture.




Figure 59 Conceptual integration of rooftop PV generation with an industrial heat pump, showing the conversion of surplus solar electricity into stored thermal energy for process heat and space heating.

5.3.3 Advanced PV Technologies and BIPV Integration

The 450 W monocrystalline module used in this study performs in the 14–20% efficiency range [60]. It suits flat and low-slope roof surfaces well, but the Stellantis facility also has sawtooth sections and vertical façade areas that conventional rack-mounted systems simply cannot reach.

Thin-film photovoltaic layers, deposited at thicknesses of 1–2 microns onto flexible substrates, can be integrated into façades, curved surfaces, and glazing elements [27]. Efficiency is lower, typically 5–7% [27], but the accessible surface area grows considerably. Within the existing

framework, what changes is mostly the performance coefficients and the spatial filtering rules; the core IFC parameterisation and Revit integration stay the same [27]. Lower efficiency per square metre can be offset by deploying across a much larger total surface, which may deepen the overall decarbonisation outcome for the site.




Figure 60 Comparison of monocrystalline PV and thin-film BIPV, highlighting expanded envelope coverage and greater decarbonisation potential.

5.3.4 Integration of Biomass for Baseload Supplementation

Solar PV provides clean electricity but remains inherently intermittent, particularly during winter and prolonged periods of low solar irradiance. Integrating biomass energy systems into industrial facilities can provide reliable baseload generation by converting locally available biomass resources, such as wood chips and agricultural residues, into thermal or electrical energy. Unlike solar PV, biomass can be dispatched on demand, improving energy security and reducing dependence on grid electricity during periods of low renewable generation.

Extending the BIM framework to include biomass boilers, fuel storage, and automated handling systems would enable comprehensive multi-vector energy planning. By incorporating biomass system parameters into the existing 5D BIM cost and performance model, decision-makers could evaluate hybrid strategies in which photovoltaic systems satisfy daytime electrical demand while biomass addresses continuous thermal loads and seasonal energy deficits. This integrated approach would support more resilient, cost-effective, and lower-carbon industrial energy systems.

References

- [1] European Commission. (2024). *Report from the Commission to the European Parliament and the Council: 2024 Carbon Market Report*. Brussels, Belgium: European Climate Action. Available at: https://climate.ec.europa.eu/eu-action/eu-emissions-trading-system-eu-ets/union-registry_en
- [2] International Renewable Energy Agency (IRENA). (2025). *Renewable Power Generation Costs in 2024*. Abu Dhabi: International Renewable Energy Agency. Available at: <https://www.irena.org/Publications>
- [3] Castellanos, S., Sunter, D. A., & Kammen, D. M. (2017). Rooftop solar photovoltaic potential in cities: How scalable are assessment approaches? *Environmental Research Letters*, 12(12), 125005. <https://doi.org/10.1088/1748-9326/aa7857>
- [4] European Committee for Standardization (CEN). (2003/2005). *Eurocode 1: Actions on structures – Part 1-3: Snow loads (EN 1991-1-3) & Part 1-4: Wind actions (EN 1991-1-4)*. Brussels, Belgium: CEN.
- [5] Woyte, A., Nijs, J., & Belmans, R. (2003). Partial shadowing of photovoltaic arrays with different system configurations: literature review and field test results. *Solar Energy*, 74(3), 217–233. [https://doi.org/10.1016/S0038-092X\(03\)00155-5](https://doi.org/10.1016/S0038-092X(03)00155-5)
- [6] MacAlpine, S., Erickson, R., & Brandemuehl, M. (2013). Characterization of power optimizer potential to increase energy yield in PV systems under partial shading conditions. *Solar Energy*, 98, 292–307. <https://doi.org/10.1016/j.solener.2013.09.031>
- [7] Batzelis, E. I., Papathanassiou, S. A., & Margaris, N. I. (2014). Explicit analytical models for PV modules/strings under partial shading conditions. *Renewable Energy*, 71, 630–638. <https://doi.org/10.1016/j.renene.2014.05.047>
- [8] Freitas, S., Catita, C., Redweik, P., & Brito, M. C. (2015). Modelling solar potential in the urban environment: State-of-the-art review. *Renewable and Sustainable Energy Reviews*, 41, 291–331. <https://doi.org/10.1016/j.rser.2014.08.060>
- [9] Eastman, C., Teicholz, P., Sacks, R., & Liston, K. (2011). *BIM Handbook: A Guide to Building Information Modeling for Owners, Managers, Designers, Engineers and Contractors* (2nd ed.). John Wiley & Sons.
- [10] Sacks, R., Eastman, C., Lee, G., & Teicholz, P. (2018). *BIM Handbook: A Guide to Building Information Modeling for Owners, Designers, Engineers, Contractors, and Facility Managers* (3rd ed.). John Wiley & Sons.
- [11] Bew, M., & Richards, M. (2008). BIM Maturity Model. *Construct IT Autumn 2008 Members Meeting*, Brighton, UK.

- [12] Succar, B. (2009). Building information modelling framework: A research and delivery foundation for industry stakeholders. *Automation in Construction*, 18(3), 357–375. <https://doi.org/10.1016/j.autcon.2008.10.003>
- [13] American Institute of Architects (AIA). (2020). *Level of Development (LOD) Specification Part I & Commentary*. BIMForum. Available at: <https://bimforum.org/lod/>
- [14] Abanda, F. H., Vidalakis, C., Oti, A. H., & Tah, J. H. M. (2015). A critical analysis of Building Information Modelling systems used in construction projects. *Advances in Engineering Software*, 90, 183–201. <https://doi.org/10.1016/j.advengsoft.2015.08.009>
- [15] Cavka, H. B., Staub-French, S., & Poirier, E. A. (2017). Developing owner information requirements for BIM-enabled project delivery and asset management. *Automation in Construction*, 83, 169–183. <https://doi.org/10.1016/j.autcon.2017.08.006>
- [16] Lu, Q., Chen, L., Lee, S., & Zhao, X. (2018). Activity theory-based analysis of BIM implementation in building O&M and first response. *Automation in Construction*, 85, 317–332. <https://doi.org/10.1016/j.autcon.2017.10.017>
- [17] Kensek, K. (2014). Integration of Environmental Sensors with BIM: case studies using Arduino, Dynamo, and the Revit API. *Informes de la Construcción*, 66(536), e044. <https://doi.org/10.3989/ic.13.151>
- [18] Preidel, C., & Borrmann, A. (2016). Towards code compliance checking on the basis of a visual programming language. *Journal of Information Technology in Construction (ITcon)*, 21, 402–421. Available at: <https://www.itcon.org/2016/25>
- [19] Asl, M. R., Zarrinmehr, S., Bergin, M., & Yan, W. (2015). BPOpt: A framework for BIM-based performance optimisation. *Energy and Buildings*, 108, 401–412. <https://doi.org/10.1016/j.enbuild.2015.09.011>
- [20] Kensek, K. (2015). Visual programming for building information modeling: energy and shading analysis case studies. *Journal of Green Building*, 10(4), 28–43. <https://doi.org/10.3992/1943-4618-10.4.28>
- [21] Crawley, D. B., Hand, J. W., Kummert, M., & Griffith, B. T. (2008). Contrasting the capabilities of building energy performance simulation programs. *Building and Environment*, 43(4), 661–673. <https://doi.org/10.1016/j.buildenv.2006.10.027>
- [22] Ramaji, I. J., & Memari, A. M. (2016). Information exchange framework for building information modeling (BIM) based structural analysis. *Journal of Computing in Civil Engineering*, 30(6), 04016035. [https://doi.org/10.1061/\(ASCE\)CP.1943-5487.0000595](https://doi.org/10.1061/(ASCE)CP.1943-5487.0000595)
- [23] Dong, B., Lam, K. P., Huang, Y. C., & Dobbs, G. M. (2007). A comparative study of the IFC and gbXML informational infrastructures for data exchange in computational design support environments. *Journal of Building Performance Simulation*, 1(1), 1–17. <https://doi.org/10.1080/19401490701836465>

- [24] Venugopal, M., Eastman, C. M., Sacks, R., & Teizer, J. (2012). Semantics of model views for information exchanges using the industry foundation class schema. *Advanced Engineering Informatics*, 26(2), 411–428. <https://doi.org/10.1016/j.aei.2012.01.005>
- [25] European Committee for Standardization (CEN). (2002). *Eurocode: Basis of structural design (EN 1990)*. Brussels, Belgium: CEN.
- [26] Crawley, D. B., Lawrie, L. K., Winkelmann, F. C., Buhl, W. F., Huang, Y. J., Pedersen, C. O., ... & Glazer, J. (2001). EnergyPlus: creating a new-generation building energy simulation program. *Energy and Buildings*, 33(4), 319–331. [https://doi.org/10.1016/S0378-7788\(00\)00114-6](https://doi.org/10.1016/S0378-7788(00)00114-6)
- [27] Kouhirostami, M., & Nourian, P. (2022). A review of current workflows and tools for Building Integrated Photovoltaics (BIPV) design. *Renewable and Sustainable Energy Reviews*, 154, 111812. <https://doi.org/10.1016/j.rser.2021.111812>
- [28] Robinson, D., & Stone, A. (2004). Solar radiation modelling in the urban context. *Solar Energy*, 77(3), 295–309. <https://doi.org/10.1016/j.solener.2004.05.010>
- [29] Stellantis. (2024). *Architectural plans and technical documentation for the San Benigno Canavese facility (Stellantis manufacturing plant)*. Internal project document.
- [30] Wilcox, S., & Marion, W. (2008). *User's manual for TMY3 data sets (No. NREL/TP-581-43156)*. National Renewable Energy Laboratory (NREL), Golden, CO. Available at: <https://www.nrel.gov/docs/fy08osti/43156.pdf>
- [31] European Committee for Standardization (CEN). (2004). *Eurocode 8: Design of structures for earthquake resistance – Part 1: General rules, seismic actions and rules for buildings (EN 1998-1)*. Brussels, Belgium: CEN.
- [32] Kamel, E., & Memari, A. M. (2019). Review of BIM's application in energy simulation: Tools, issues, and solutions. *Automation in Construction*, 97, 164–180. <https://doi.org/10.1016/j.autcon.2018.11.005>
- [33] Cui, Y., et al. (2021). A review of visual programming in AEC industry. *Automation in Construction*, 128, 103756. <https://doi.org/10.1016/j.autcon.2021.103756>
- [34] Lu, W., et al. (2013). Building information modeling and changing construction practices. *Journal of Construction Engineering and Management*, 139(1), 108–116. [https://doi.org/10.1061/\(ASCE\)CO.1943-7862.0000569](https://doi.org/10.1061/(ASCE)CO.1943-7862.0000569)
- [35] Sanjuán-Micó, A., et al. (2020). BIM-to-BEM workflow for energy performance assessment. *Applied Sciences*, 10(5), 1752. <https://doi.org/10.3390/app10051752>
- [36] Gledson, B. J., & Greenwood, D. (2016). Surveying the extent and use of 4D BIM in the UK. *Journal of Information Technology in Construction (ITcon)*, 21, 57–71. Available at: <https://www.itcon.org/2016/4>

- [37] Smith, P. (2014). BIM & the 5D Project Cost Manager. *Procedia - Social and Behavioral Sciences*, 119, 475–484. <https://doi.org/10.1016/j.sbspro.2014.03.053>
- [38] Boje, C., et al. (2020). Towards a semantic Construction Digital Twin: Directions for future research. *Automation in Construction*, 114, 103179. <https://doi.org/10.1016/j.autcon.2020.103179>
- [39] Boton, C., et al. (2015). 4D BIM based simulation of construction logistics. *Procedia Engineering*, 123, 42–49. <https://doi.org/10.1016/j.proeng.2015.10.055>
- [40] Moon, H., et al. (2014). Application of 4D CAD in construction logistics. *Automation in Construction*, 44, 44–52. <https://doi.org/10.1016/j.autcon.2014.03.023>
- [41] Davies, R., & Harty, C. (2013). Implementing 'Site BIM': a case study of ICT innovation on a large hospital project. *Automation in Construction*, 30, 15–24. <https://doi.org/10.1016/j.autcon.2012.11.024>
- [42] Chau, K. W., et al. (2004). 4D dynamics management system for construction planning and resource utilization. *Advances in Engineering Software*, 35(3-4), 239–248. <https://doi.org/10.1016/j.advengsoft.2004.03.003>
- [43] Lu, K., et al. (2016). Application of BIM 5D in the cost management of construction projects. *Journal of Management in Engineering*, 32(5), 04016009. [https://doi.org/10.1061/\(ASCE\)ME.1943-5479.0000450](https://doi.org/10.1061/(ASCE)ME.1943-5479.0000450)
- [44] Vigneault, A., et al. (2020). Framework for integrating life-cycle costing in BIM. *Journal of Building Engineering*, 28, 101032. <https://doi.org/10.1016/j.jobbe.2019.101032>
- [45] Kim, K., et al. (2020). An integrated 5D BIM approach to cost and schedule management. *Automation in Construction*, 118, 103254. <https://doi.org/10.1016/j.autcon.2020.103254>
- [46] Mitchell, D. (2012). *5D BIM: Creating cost certainty and better buildings*. RICS Publishing.
- [47] Lu, Q., et al. (2020). Developing a Digital Twin at Building and City Levels: Case Study of West Cambridge Campus. *Journal of Management in Engineering*, 36(3), 05020004. [https://doi.org/10.1061/\(ASCE\)ME.1943-5479.0000763](https://doi.org/10.1061/(ASCE)ME.1943-5479.0000763)
- [48] Khajavi, S. H., et al. (2019). Digital Twin: Vision, benefits, boundaries, and creation for buildings. *IEEE Access*, 7, 147406–147419. <https://doi.org/10.1109/ACCESS.2019.2946515>
- [49] Andriamamonjy, A., et al. (2018). An automated BIM-based process for generating BEM models. *Automation in Construction*, 91, 39–51. <https://doi.org/10.1016/j.autcon.2018.03.004>
- [50] Gerrish, T., et al. (2017). Using BIM capabilities to improve existing building energy modelling practices. *Energy and Buildings*, 138, 544–558. <https://doi.org/10.1016/j.enbuild.2016.12.083>

- [51] Parn, E. A., Edwards, D. J., & Sing, M. C. (2017). The building information modelling trajectory in facilities management: A review. *Automation in Construction*, 75, 45–55. <https://doi.org/10.1016/j.autcon.2016.12.003>
- [52] Bidram, A., Davoudi, A., & Balog, R. S. (2012). Control and circuit techniques to mitigate partial shading effects in photovoltaic arrays. *IEEE Journal of Photovoltaics*, 2(4), 532–546. <https://doi.org/10.1109/JPHOTOV.2012.2202879>
- [53] Silvestre, S., & Castañer, L. (2007). Forward-bias performance of faulty solar cells. *Progress in Photovoltaics: Research and Applications*, 15(8), 653–662. <https://doi.org/10.1002/pip.759>
- [54] Petrone, G., Spagnuolo, G., & Vitelli, M. (2007). Analytical model of mismatched photovoltaic fields by means of Lambert W-function. *Solar Energy Materials and Solar Cells*, 81(12), 2659–2675. <https://doi.org/10.1016/j.solmat.2007.07.009>
- [55] Sampaio, A. Z., & Simões, D. (2016). Maintenance of buildings using BIM methodology. *Journal of Information Technology in Construction (ITcon)*, 21, 303–318. Available at: <https://www.itcon.org/2016/20>
- [56] Jalaei, F., & Jrade, A. (2015). Integrating building information modeling (BIM) and LEED system at the conceptual design stage of sustainable buildings. *Sustainable Cities and Society*, 18, 95–107. <https://doi.org/10.1016/j.scs.2015.06.007>
- [57] Regione Piemonte. (2024). *Prezziario Regionale delle Opere Pubbliche*. Turin, Italy. Available at: <https://www.regione.piemonte.it/web/temi/sviluppo/opere-pubbliche/prezzario-regionale>
- [58] Roudsari, M. S., & Pak, M. (2013). Ladybug: a parametric environmental plugin for Grasshopper to help designers create an environmentally-conscious design. In *Proceedings of the 13th international conference of the IBPSA* (pp. 3128–3135).
- [59] European Commission. (2023). *Photovoltaic Geographical Information System (PVGIS)*. Joint Research Centre (JRC). Available at: https://joint-research-centre.ec.europa.eu/pvgis-photovoltaic-geographical-information-system_en
- [60] International Energy Agency (IEA). (2024). *Trends in Photovoltaic Applications*. IEA Photovoltaic Power Systems Programme (PVPS). Available at: <https://iea-pvps.org/>
- [61] Autodesk. (2024). *Revit API Developers Guide*. Autodesk, Inc. Available at: <https://help.autodesk.com/view/RVT/2024/ENU/>
- [62] McNeel, R. (2024). *Rhinoceros 3D*. Robert McNeel & Associates. Available at: <https://www.rhino3d.com/>
- [63] Mermoud, A., & Wittmer, B. (2014). *PVsyst user's manual*. PVsyst SA, Switzerland. Available at: <https://www.pvsyst.com/>

[64] Marion, W., & Urban, K. (1995). *User's manual for TMY2s: Typical Meteorological Years*. National Renewable Energy Laboratory (NREL), Golden, CO. Available at: <https://www.nrel.gov/docs/legosti/old/7356.pdf>

[65] ISPRA. (2023). *Italian Greenhouse Gas Inventory 1990-2021: National Inventory Report 2023*. Italian Institute for Environmental Protection and Research (ISPRA), Rome, Italy. Available at: <https://www.isprambiente.gov.it/>

[66] Ballarini, I., & Corrado, V. (2024). *Renewable energy sources (RES) and related technologies* [Lecture notes]. Course: Energy performance design and indoor environmental quality. Department of Energy, Politecnico di Torino.