A MICROSERVICE-ORIENTED INTEGRATION LAYER FOR LV MANAGEMENT TOOLS

Gil Sampaio[1], Jacinta Ferreira[1], José Sousa[1], João Fernandes[2], Raquel Figueiredo[2], José Oliveira[2]*

[1]INESC TEC, Porto, Portugal
[2] Eneida.io, Coimbra, Portugal
**gil.s.sampaio@inesctec.pt*

Keywords: LV GRIDS, MICROSERVICES, INTEROPERABILITY, STATE ESTIMATION, OPERATING ENVELOPES

Abstract

With the growing need to manage increasingly complex and decentralised low voltage (LV) grids, Distribution System Operators (DSOs) are increasingly interested in interoperable platforms that can integrate advanced analytics and control functions from different vendors. In the ATE initiative (NextGenerationEU), INESC TEC and Eneida combine LV monitoring and management with a microservice-oriented integration approach. This paper describes how the eneida DeepGrid® platform was extended to integrate three INESC TEC microservices: Data-driven State Estimation (DdSE), Voltage Control (DdVC), and Dynamic Operating Envelopes (DdOE), using the eneida DeepGrid® Software Development Kit (SDK) as the common integration mechanism. A composite microservice, the LV Manager, encapsulates and orchestrates the interactions between the DdSE, the DdVC, and the DdOE, while exposing a single interface to the platform. Quantitative results are reported from a pre-deployment validation on a representative LV test grid using real consumption and PV generation profiles; in parallel, a demonstration across 15 urban secondary substations is underway and field data is being gathered for follow-up assessment.

1 Introduction

LV grids are evolving towards higher complexity due to the continuous growth of decentralised resources and the increasing need for situational awareness at the secondary substation and feeder levels [1], [2]. This evolution requires a fundamental move away from passive network management toward active, data-driven control. In this context, DSOs no longer seek rigid, monolithic systems that require complete replacement for every upgrade. Instead, the modern requirement is for agile software platforms that support incremental feature deployment, allowing functions to be integrated as "plug-and-play" add-ons to existing infrastructure.

Despite the clear demand, the transition from research to operation often hits an "integration wall". While many advanced LV tools have been successfully validated as robust standalone applications [3], [4], their real-world adoption is frequently stifled by bespoke interfaces, redundant data handling, and tight coupling between components. This paper addresses that barrier by proposing a microservice-oriented integration layer that allows LV management tools from different providers to be connected through a common interface and shared data pipelines, reducing integration effort and limiting vendor lock-in.

The key contributions of this work are listed below:

1. A microservice-oriented integration approach for deploying third-party LV analytics and control functions in eneida DeepGrid® via the eneida DeepGrid® SDK.
2. The LV Manager composite microservice, which bundles DdSE, DdVC and DdOE behind a single interface.
3. Validation on a representative LV test grid using real consumption/generation profiles demonstrating end-to-end consistency and reporting application-level KPIs, while the field demonstration is ongoing.

The remainder of the paper is organised as follows. Section 2 describes the microservice-oriented integration layer and the LV Manager orchestration design. Section 3 presents the proof-of-concept deployment and the obtained results. Section 4 concludes the paper and outlines directions for extending the evaluation, including quantitative runtime and integration KPIs and behaviour under missing-data conditions.

2 Methodology

This section details the microservice-oriented integration approach used to incorporate advanced LV analytics and control capabilities into the platform without requiring bespoke integrations for each tool. The methodology is founded on three pillars: (i) a standard integration mechanism via the eneida DeepGrid® SDK, (ii) a shared external data backbone, and (iii) a composite microservice pattern used to orchestrate sub-services.

2.1 eneida DeepGrid® SDK integration layer

The eneida DeepGrid® SDK abstracts the complexity of microservice interaction by exposing a standard interface for integration alongside native applications. It is delivered with a containerized development environment to ensure consistent runtimes and reduce setup time for third-party developers.

2.1.1 Overview and components: Eneida's platform architecture (Figure 1) utilizes stable data models and standardized interfaces to shield external applications from internal platform evolution. This is achieved through two complementary SDKs:

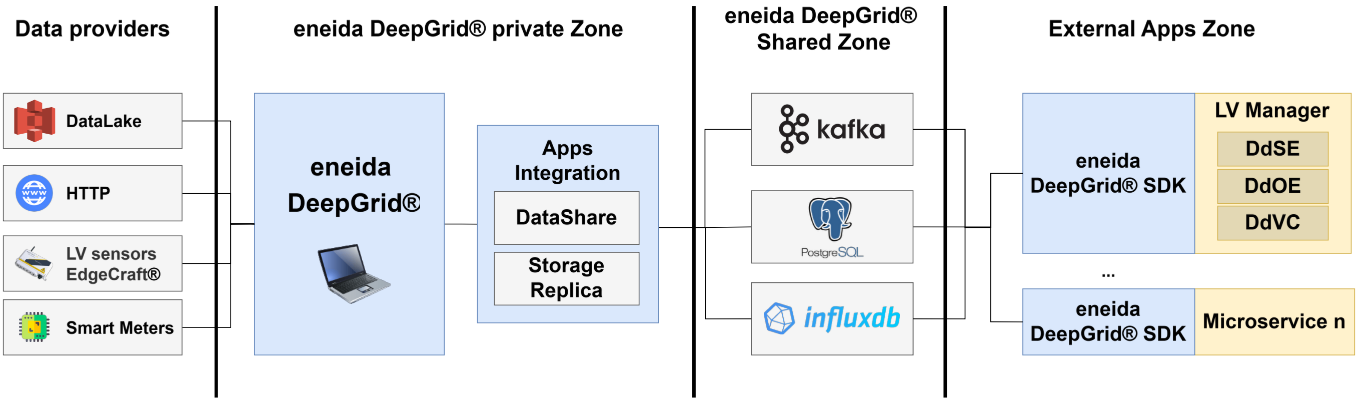

Figure 1: Architecture of the eneida DeepGrid® SDK and its external data pipelines (Kafka, InfluxDB, PostgreSQL), exposing a standard interface for integrating third-party analytics microservices alongside native applications.

- eneida EdgeCraft®: Operates at the device and edge level, handling local intelligence, data acquisition, and raw data ingestion to isolate the cloud platform from device management complexities.
- eneida DeepGrid® SDK: Provides a modular, technology-agnostic framework for analytics and control micro-services at the cloud level.

Together, the two SDKs enforce a clear separation of concerns between edge and cloud, while keeping a coherent end-to-end data flow, even when different providers are contributing at edge and cloud-levels. This layered design allows each layer to evolve independently, and the eneida DeepGrid® SDK is at the center as the main integration contract for external cloud-based analytics and control microservices, providing a stable, extensible, and provider-independent framework for flexible and interoperable LV grid apps.

2.1.2 Integration architecture and data interface: The integration architecture of the eneida DeepGrid® SDK is built on three well-established external technologies, promoting interoperability and avoiding technology lock-in: (i) Apache Kafka provides a scalable and resilient real-time streaming mechanism for both data consumption and publication between the platform and external microservices; (ii) InfluxDB enables access to historical time-series measurements required for analytics, monitoring and data-driven estimation; (iii) PostgreSQL stores and provides relational data including grid models, network topology, configuration parameters, metadata, and reference information required by analytics and control microservices.

The SDK wraps these into a unified programming interface, allowing developers to integrate new microservices without interacting directly with the underlying pipelines. This abstraction reduces the burden of connectivity and baseline data handling; tool-specific validation and domain checks remain within each microservice. Through this interface, apps can: subscribe to real-time data streams produced by edge-devices and to enriched data from eneida DeepGrid® internal analytics; publish processed results, calculated indicators and control signals back to the platform; query historical time-series data, including raw measurements and enriched datasets; access relational data such as grid models, parameters, metadata and contextual information.

2.1.3 Development environment and scalability: A platform's ability to grow depends largely on how quickly and reliably new services can be brought into it. To support this, the eneida DeepGrid® SDK is delivered with a fully prepared containerized development environment that includes all required runtime components, libraries and configuration files, removing infrastructure setup as a barrier to the integration. In addition to local development support, the SDK is complemented by: (i) device simulators to allow partners to validate application logic without requiring access to physical hardware; (ii) a shared cloud-based development environment provided by the platform, which allows partners to develop and validate their apps under conditions that closely resemble operational deployments, while keeping isolation from production environments.

The eneida DeepGrid® SDK was designed to keep an ecosystem open to a wide range of contributors. The initial implementation is in Python, this decision was taken for its suitability in data-driven analytics, rapid prototyping and integration with streaming and database technologies. However, the data models, interfaces and integration patterns are defined in a language-agnostic manner, so that equivalent SDK implementations in other languages can be developed in future phases without the need to redesign the architecture.

2.2 LV Manager as a composite microservice

To avoid integrating multiple function-specific services independently into the platform, we implement LV Manager as a composite microservice. From the perspective of eneida DeepGrid®, LV Manager behaves as a single external service accessed through a single eneida DeepGrid® SDK connection. Internally, LV Manager encapsulates three tools and manages their ordering, inputs, and data handoff, enforcing the following orchestration chain:

- DdSE combines historical data with available partial system information to produce an updated estimate of the LV operating state.
- DdVC uses the estimated state together with flexibility bids (when available) to compute active-power adjustments that mitigate voltage-limit violations, producing a corrected operating point and the corresponding control actions to flexible resources.

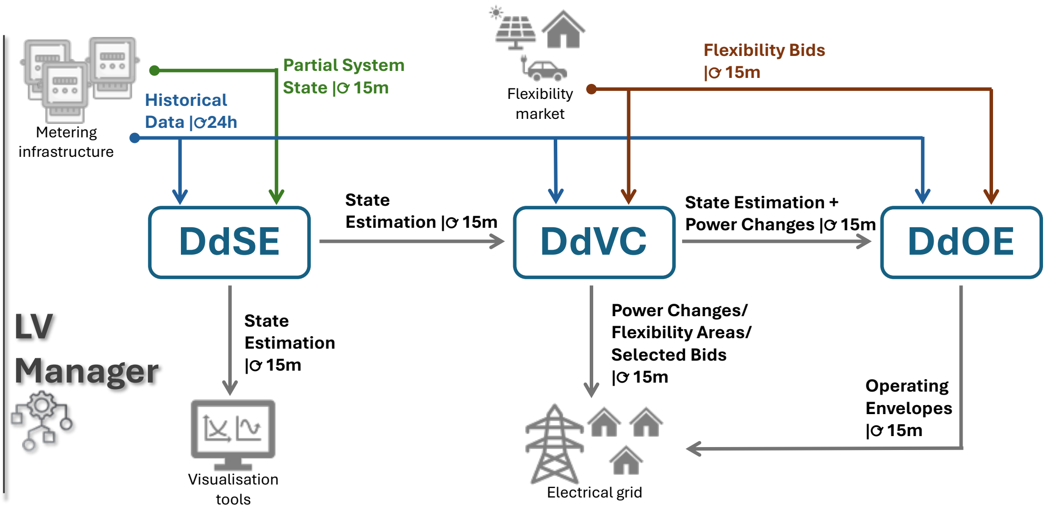

Figure 2: Conceptual view of the LV Manager as a composite microservice, showing how it orchestrates the DdSE, DdVC and DdOE sub-services and manages the data flows between them.

- DdOE computes admissible operating bounds around the operating point, translating network feasibility into actionable flexibility limits (envelopes) that can be used to constrain or validate dispatch requests.

In operational terms, LV Manager executes the chain at a configurable cadence. For each cycle, inputs are aligned to a consistent snapshot so that state estimation, control, and envelope computation are based on the same data context despite heterogeneous refresh rates. Intermediate outputs are explicitly passed between stages: DdVC operates on the state estimate produced by DdSE, and DdOE uses the operating point resulting from the upstream stages as its reference.

Figure 2 summarises the resulting information flow. Two inputs feed DdSE: (i) historical data (blue path) and (ii) partial system information (green path). The state estimate is then forwarded to DdVC, which also ingests flexibility bids (orange path), when available, to compute control actions to flexible assets (downstream actuation path). Finally, DdOE uses the upstream context and the flexibility information to compute operating envelopes, which are returned as admissible bounds that can be used to constrain flexibility usage and support operational decision-making.

2.3 Tool operational assumptions and data requirements

The workflow operates under a mixed-observability LV data regime, combining historical smart-meter datasets with a subset of near real-time measurements. In the demonstration, near real-time meters are available at a 15-minute cadence, while the remaining meters are updated daily; the near real-time subset represents around 10% of the population. Historical data supports model calibration and periodic refresh (including sensitivity updates), while the near real-time subset anchors each execution cycle to the current operating conditions.

Two operation modes are supported. In online operation, the chain is executed periodically using the latest aligned measurements and available operational inputs (including flexibility bids when available). In preventive operation, the same chain is executed over a selected horizon using forecasted operating conditions, enabling constraint-aware planning of flexibility usage.

2.4. Tool algorithmic overview

2.4.1 DdSE [6]: targets near real-time reconstruction of LV voltage states under partial observability, i.e., when only a subset of smart meters provides timely measurements. The estimator is data-driven and is designed to operate without requiring full knowledge of LV grid parameters. A key implementation follows an analog-search paradigm: historical smart-meter datasets are used to build a library of past operating states, and, at runtime, the current partially observed snapshot is matched against this library to infer unobserved voltages through similarity-based reconstruction. In operational terms, DdSE supports mixed refresh regimes by combining: (i) historical data to maintain a representative library of operating conditions and seasonal patterns, and (ii) near real-time measurements (from the subset of meters/sensors with faster communication) to anchor the current snapshot. Optional exogenous/contextual variables (e.g., calendar effects, weather proxies) can be incorporated to improve matching robustness.

2.4.2 DdVC [7]: computes active-power adjustments to mitigate LV voltage-limit violations using a data-driven linearisation of voltage responses to power variations. The method relies on voltage-to-active-power sensitivity factors learned from measurements (e.g., regression-based estimation of $\Delta V/\Delta P$ relationships), enabling fast evaluation of voltage impacts without explicit full-parameter models. Control is formulated as a linear optimisation over active power adjustments ΔP of available flexible resources, subject to (i) resource bounds (offered flexibility ranges) and (ii) voltage constraints expressed through the learned sensitivities. If available flexibility is insufficient to fully restore feasibility, DdVC can quantify residual corrective needs at the secondary substation (e.g., transformer voltage set-point support) as a complementary action channel. A key operational feature is that sensitivity learning can be made more robust and deployable via (i) privacy-preserving estimation mechanisms and (ii) operating-point dependent ("varying-coefficient") sensitivities, improving control accuracy under changing loading and voltage conditions. In addition, DdVC can be embedded in coordinated flexibility workflows (e.g., market-based or bid-based flexibility activation), where bids provide the feasible adjustment ranges.

2.4.3 DdOE [7]: computes time-varying admissible active-power bounds (operating envelopes) per customer connection point or node, such that operating within these bounds preserves network feasibility with respect to voltage limits. Conceptually, DdOE turns network constraints into an actionable constraint interface that can be consumed by flexibility platforms and customer/aggregator control, complementing corrective voltage control with preventive admissibility limits. Algorithmically, DdOE is implemented through sensitivity-based optimisation, leveraging the same linearised voltage model used for DdVC. Envelope computation can be posed as a set of linear programmes that maximise admissible injection and/or consumption while enforcing voltage constraints. Beyond individual envelopes, the same formulation supports aggregated ones at substation level when a feeder-wide admissibility constraint is required.

3 Results

The project includes an ongoing demonstration across 15 urban secondary substations integrated in eneida DeepGrid®, where the LV Manager composite microservice is deployed to operate on field measurements. At the time of writing, field data from the demonstration is still being gathered and consolidated to support statistically representative performance reporting. Therefore, the quantitative results presented in this section focus on a pre-deployment validation performed on a representative simulation environment built upon the 33-node LV network shown in Figure 3, using load data from a trial by the CER in Ireland [8] and PV generation data from the SuSTAINABLE project [9].

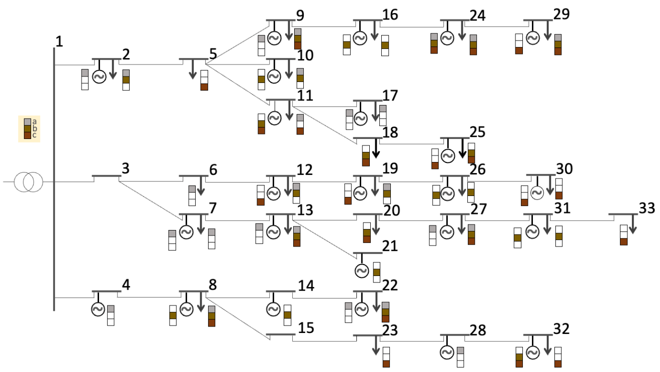

Figure 3: 33-node LV grid used for testing the methodology.

All three tools were calibrated using one year of historical data at 30-minute resolution. The intent is to operate a realistic LV "data regime": (i) a subset of measurements is available at an operational cadence suitable for near real-time execution, while (ii) the remaining inputs contribute primarily through historical datasets that support model fitting, sensitivity updates, and envelope computation. This reflects the common DSO reality in LV, where observability increases progressively and data timeliness is heterogeneous.

DdSE: The accuracy of the DdSE component is vital as it provides the foundational state for subsequent control actions. Running the tool for a 30-day horizon yielded a Maximum Absolute Deviation (MAD) of 5.14 V, a Mean Absolute Error (MAE) of 0.74 V, and a Mean Absolute Percentage Error (MAPE) of 0.31%. The low magnitude of these errors indicate that the data-driven model provides highly reliable and accurate results for LV monitoring. Figure 4 illustrates the comparison between real and estimated voltages over a 48-hour period. The estimated profile accurately captures the daily variations and peak fluctuations with only minor deviations, maintaining consistency with the reported metrics.

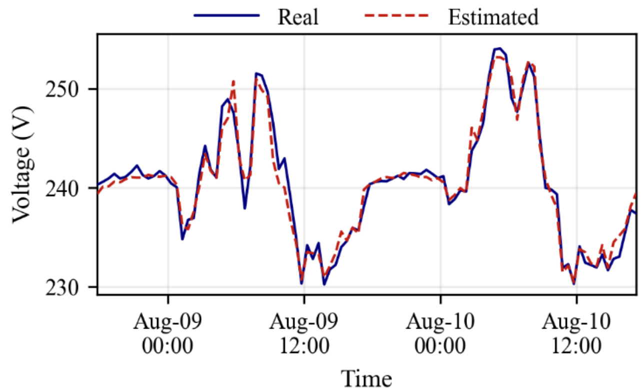

Figure 4: Measured ("Real") vs estimated voltage profile for a representative client over a 48-hour period.

DdVC: The performance of the DdVC tool can be assessed through the Violation Frequency Reduction (VFR). During the 30-day testing period, a VFR of 94.17% was achieved, indicating a substantial decrease in the occurrence of voltage violations after the control strategy was applied. This result demonstrates the strong effectiveness of the tool in maintaining voltage levels within acceptable bounds. Figure 5 presents the comparison between the initially expected voltage profile and the voltage profile after the control action, for one client, over a 48-hour period.

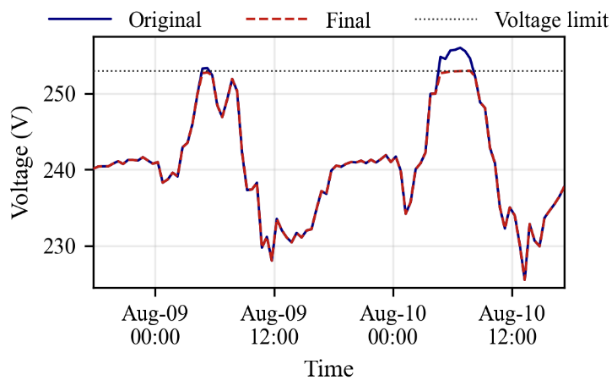

Figure 5: Voltage profile before ("Original") and after ("Final") the DdVC control action for a representative client over a 48-hour period, including the maximum voltage limit.

A voltage violation of the maximum limit was observed on both days, with a more pronounced deviation on the second day. After applying the selected power adjustments, the voltage profile was reduced to at or below the limit.

DdOE: The performance of the DdOE tool can be assessed through the Additional Dispatchable Power (ADP), expressed as a percentage of the admissible dispatch range. Over a 30-day testing period, the envelopes enabled an ADP of 76.79%, meaning that, starting from the corrected operating point, an additional 76.79% of the admissible dispatchable flexibility could be safely activated while remaining within the network constraints. This result demonstrates the practical value of DdOE in translating network feasibility into an actionable flexibility margin for operation.

Figure 6 illustrates the DdOE output for a representative timestamp. For clarity, a symmetric ±10 kW range is shown; in practice, limits are updated by the tool and may vary across nodes and time. The plot shows the node-level feasible envelopes and the DdVC operating point ("Original"), which is corrected and therefore always lies within the DdOE envelope.

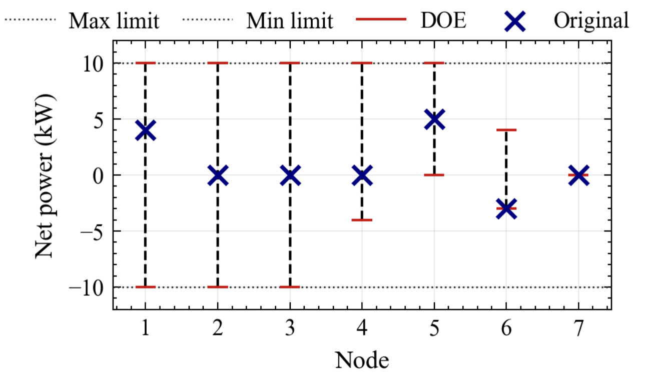

Figure 6: Illustrative snapshot of DdOE node-level feasible power envelopes at a representative timestamp (±10 kW reference range), with the corrected operating point from DdVC ("Original").

4 Conclusion

This paper presented a microservice-oriented integration approach for deploying LV analytics and control functions in eneida DeepGrid® via the eneida DeepGrid® SDK, avoiding bespoke interfaces and tight coupling. Building on this integration layer, we introduced LV Manager as a composite microservice that encapsulates and orchestrates a multi-stage LV workflow behind a single platform-facing interface. The chain combines DdSE, DdVC, and DdOE, explicitly managing their data dependencies so that state reconstruction, corrective control actions, and envelope computation remain consistent within each execution cycle.

Quantitative results were obtained from a pre-deployment validation on a representative 33-node LV test grid driven by real consumption and PV generation profiles, demonstrating end-to-end workflow coherence and application-level KPIs for the three tools. In parallel, a demonstration across 15 urban secondary substations integrated in eneida DeepGrid® is underway, and field data is currently being gathered for follow-up assessment under operational conditions.

Future work will complete the field-data evaluation and add platform-level KPIs, including end-to-end latency, robustness to missing or delayed measurements, and behaviour under partial service failures, supporting deployment at larger scale and under changing operating conditions.

5 Acknowledgements

This work is co-financed by Component 5 - Capitalization and Business Innovation, integrated in the Resilience Dimension of the Recovery and Resilience Plan within the scope of the Recovery and Resilience Mechanism (MRR) of the European Union (EU), framed in the Next Generation EU, for the period 2021 - 2026, within project ATE, with reference 56.

6 References

[1] R. Tonkoski, D. Turcotte, T.H.M. El-Fouly: 'Impact of high PV penetration on voltage profiles in residential neighborhoods', *IEEE Trans. Sustain. Energy*, 2012.
[2] D. Lima, G. Sampaio: 'Topology reconstruction of low voltage grids using genetic algorithms', in *Proc. IEEE Int. Conf. on Systems, Man, and Cybernetics (SMC)*, 2025.
[3] D.F. Lopes, M. Simões, G. Sampaio, *et al.*: 'From home energy management system local flexibility to low-voltage predictive grid management', *CIRED*, 2020.
[4] G. Sampaio, C. Gouveia, R. Bessa, *et al.*: 'EUniversal's smart grid solutions for the coordinated operation and planning of MV and LV networks with high EV integration', in *CIRED Porto Workshop 2022: E-mobility and power distribution systems*.
[5] M. Simões, G. Sampaio, C. Madureira, *et al.*: 'Avoid technical problems in LV networks: from data-driven monitoring to predictive control', in *Proc. CIRED*, 2019.
[6] R.J. Bessa, G. Sampaio, V. Miranda, J. Pereira: 'Probabilistic low-voltage state estimation using analog-search techniques', in *Proc. Power Systems Computation Conference (PSCC)*, 2018.
[7] G. Sampaio, R.J. Bessa, C. Gonçalves, C. Gouveia: 'Conditional parametric model for sensitivity factors in LV grids: a privacy-preserving approach', in *Proc. Power Systems Computation Conference (PSCC)*, 2022.
[8] Commission for Energy Regulation (CER): 'Electricity smart metering customer behaviour trials findings report', Dublin, Ireland, Tech. Rep. CER/11/080a, 2011.
[9] R.J. Bessa, A. Trindade, V. Miranda: 'Spatial–temporal solar power forecasting for smart grids', *IEEE Trans. Ind. Informatics*, 2015, 11, (1), pp. 232–241.