Materials Characterization 240 (2026) 116767

Contents lists available at ScienceDirect

Materials Characterization

journal homepage: www.elsevier.com/locate/matchar

SPARSE — Efficient high-resolution SEM imaging of rare microstructural features across large areas by selective rescanning

Tom Reclik [a,*], Jan Gerlach [b], Maximilian A. Wollenweber [a], Yannis P. Korkolis [b], Sandra Korte-Kerzel [a], Ulrich Kerzel [a]

[a] Institute of Physical Metallurgy and Materials Physics, RWTH Aachen University, Aachen, 52074, Germany
[b] Institute of Forming Technology and Lightweight Components, TU Dortmund, Dortmund, 44227, Germany

ARTICLE INFO

Keywords:
Scanning electron microscopy
Large-area imaging
Rare feature
Dual-phase steel
Microstructural damage

ABSTRACT

Characterization of rare microstructural features in scanning electron microscopy (SEM) requires imaging large areas at high resolution. This leads to prohibitively long acquisition times. We present an open-source Python framework that addresses this bottleneck through a two-stage approach: a fast scan identifies regions of interest, which are then selectively rescanned with imaging parameters suitable for quantitative analysis. The framework defines a generic microscope interface and a modular detection interface, allowing adaptation to different microscope platforms and detection methods. Scanning, detection, and rescanning are parallelized using separate processes, ensuring that computation time does not extend acquisition time. The two processes communicate exclusively through queues, avoiding shared mutable state and eliminating the need for explicit synchronization. We validate the framework on damage detection in dual-phase DP800 steel using a Tescan Clara SEM. For a representative configuration a detection rate of 99 % is achieved at approximately 58 % of the conventional acquisition time. At 95 % detection rate, acquisition time drops to 19 %. These estimates are based on the ratio of scanned pixels; the achievable savings depend on the detection method, material, feature type, and the instrument. The complete implementation is available at https://doi.org/10.5281/zenodo.20321819.

1. Introduction

Characterizing rare microstructural features such as damage sites in dual-phase steels [1,2] or non-metallic inclusions in clean steels [3] requires surveying large sample areas at high spatial resolution. Conventional approaches to this problem, whether manual inspection or automated panoramic imaging [2], treat every part of the sample equally, acquiring high-resolution data across the entire region of interest. Yet for rare features, such as damage in dual-phase steels or inclusions in 16MnCrS5 case-hardening steel [4], the vast majority of the imaged area contains no relevant information. This creates a fundamental inefficiency: acquisition time scales with the total surveyed area rather than with the amount of relevant information, making comprehensive studies prohibitively slow.

Automated image acquisition in the SEM is an established and active field. Large-area characterization is commonly achieved by acquiring many overlapping high-resolution tiles and stitching them into a single panorama [2,5,6]. Throughput can also be increased at the hardware level: multi-beam SEMs acquire many fields simultaneously using arrays of tens of beams, enabling high-throughput imaging of large areas, as applied in connectomics [7,8]. Automated acquisition further underpins dedicated characterization techniques, such as in-situ mechanical testing in the SEM, where digital image correlation combined with EBSD maps the evolution of microstructure and strain during deformation in dual-phase steel [9,10], and automated particle analysis, in which feature detection is coupled with EDS to characterize non-metallic inclusions in steel, an established workflow that is also implemented in commercial systems [11]. Most of these approaches acquire data uniformly across the surveyed area, or rely on specialized hardware to accelerate that acquisition.

Selective acquisition strategies reduce this cost by concentrating high-resolution imaging on the relevant locations rather than the full area. Several approaches have been demonstrated. Merryman and Kovačević used an adaptive multirate scheme for fluorescence microscopy, acquiring a fast low-resolution survey and then sampling identified regions at higher rate [12]. In EBSD, Tong et al. accelerated orientation mapping by first recording a forescatter electron image and then collecting electron backscatter diffraction patterns only at a sparse set of points guided by that image, reconstructing the full map from this reduced sampling [13]. A more aggressive class of

* Corresponding author.
E-mail address: reclik@imm.rwth-aachen.de (T. Reclik).

https://doi.org/10.1016/j.matchar.2026.116767
Received 22 April 2026; Received in revised form 2 July 2026; Accepted 11 July 2026
Available online 24 July 2026

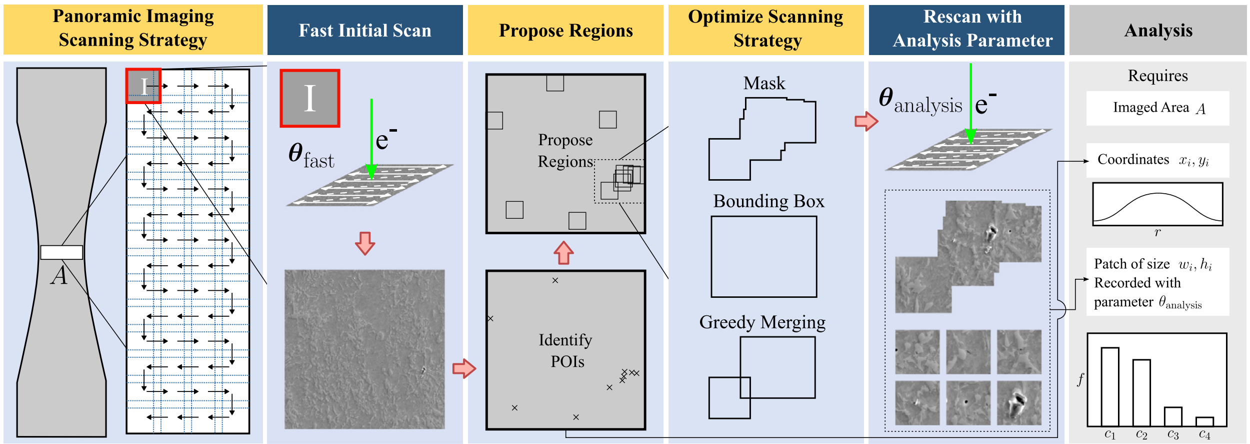

Fig. 1. Overview of the automated scanning framework SPARSE. From left to right: the panoramic imaging strategy defines the tile grid covering the target area A. Each tile is first acquired with fast imaging parameters θ_{fast}. From the fast scan, POIs are identified and regions are proposed around detected candidates. The proposed regions are then optimized into a scanning strategy. Two approaches are available depending on microscope capabilities: bounding box extraction for systems limited to rectangular scans, or greedy merging of a binary mask for systems supporting arbitrary scan regions. Selected regions are rescanned with analysis parameters θ_{analysis}. Finally, the analysis receives the coordinates of each rescan patch relative to the panorama and the corresponding high-resolution images. Colour coding of the top row indicates operations performed by the microscope (blue) versus the framework software (yellow), with downstream analysis shown in grey.

methods chooses each measurement location sequentially during acquisition, based on all previous measurements [14,15]; such dynamic sampling can minimize the number of measurements but requires a tight measure-decide-actuate loop. Within the SEM, Dahmen et al. showed that adaptive pixel-dwell sampling based on initial low-dose images substantially reduces acquisition time [16], and Clusiau et al. demonstrated feature tracking to image only regions of interest at high resolution [17]. These approaches demonstrate the potential of selective imaging strategies, but differ in how candidates are identified and how the imaging resources are allocated. When the initial scan is performed at reduced quality to accelerate acquisition, a trade-off arises between detection completeness and efficiency: more aggressive parameter reduction increases time savings but risks missing small or low-contrast features. Existing approaches have not systematically characterized this trade-off, and most implementations are tightly coupled to specific hardware or detection methods, limiting their transferability across platforms and applications.

In this work, we present an open-source framework for efficient large-area SEM characterization of rare features (SPARSE — Selective Parallelized Adaptive Rescanning for SEM Efficiency) illustrated in Fig. 1. The framework is built on two principles: modularity and parallelization. Detection methods and microscope control are implemented as interchangeable components; users implement a microscope interface class providing the required methods for their specific hardware, and a detection function appropriate for their application. This separates the open-source scanning logic from vendor-specific microscope control. Scanning, detection, and rescanning are parallelized to maximize microscope utilization, eliminating idle time during image analysis. The framework is implemented as a Python library, allowing integration into existing Python-based analysis workflows. Users define their scanning task through a configuration specifying the survey area, resolution tiers, and detection parameters. Starting from provided examples, the default scripts cover common tasks such as surveying large rectangular regions. The primary user decision lies in configuring the detection pipeline: how aggressively to reduce acquisition parameters, such as dwell time or resolution, for the initial fast scan, and which detection method to apply to the resulting images. For parametric detection methods, a sweep mode allows exploration of the accuracy–efficiency trade-off before committing to a full acquisition. An example implementation for TESCAN instruments via SharkSEM is provided.

The framework's performance is validated through a case study on microstructural damage characterization in dual-phase steel. This is a well-investigated research area that leverages statistical information on rare, heterogeneously distributed microstructural voids to increase component performance [6,18–20]. In the present case, we apply the framework to a dual-phase steel, where damage sites typically require both high-resolution and large area imaging for accurate quantification [4,21,22]. We demonstrate substantial time savings compared to conventional previously applied [2,6] full-area high-resolution imaging and characterize the trade-off between detection completeness and acquisition efficiency. The complete implementation is available as open-source software.

2. Framework architecture

The following subsections describe the terminology (Section 2.1), design principles (Section 2.2), user workflow (Section 2.3), scanning procedure (Section 2.4), and implementation details (Section 2.5).

2.1. Nomenclature

Throughout this work, we use the following terminology:

- *Panorama*: The complete tiled area on the sample to be scanned, covering an area too large for manual image-by-image acquisition.
- *Tile*: A single field of view within the panorama.
- *Partition*: A subdivision of each tile into equally sized sections to enable parallel workflows. The number of partitions is configurable; in this work, each tile is divided into four parts.
- *POI*: Point of interest; a candidate location that may contain a feature. Whether a POI corresponds to an actual feature of interest or to an artifact is determined by the subsequent analysis and is independent of the framework.
- *Region*: Region of interest; a region around one or more POIs to be rescanned at higher quality.
- θ_{fast}: Vector of imaging parameters for the fast initial scan, such as resolution, dwell time, working distance, acceleration voltage, and beam current. These are chosen to accelerate acquisition while retaining sufficient image quality for POI identification.
- $\theta_{analysis}$: Vector of imaging parameters required for subsequent quantitative analysis. These define the target image quality and are determined by the requirements of the downstream analysis.

2.2. Design principles

The framework is built on two core principles: Modularity and Parallelization.

Modularity. The framework separates microscope control from the scanning logic through a generic microscope interface. Users provide an implementation of this interface for their specific hardware, allowing the framework to be adapted to different microscope models. Similarly, detection algorithms are interchangeable components: users implement a detection function through the provided interface, choosing an approach appropriate for their application. An example implementation using thresholding and clustering is provided with the framework. Additional functionality for large-area scanning — such as failure recovery, multi-session scanning, and parameter drift compensation — is available as optional modules, which can be further expanded. The specific architectural choices and their rationale are described in Section 2.5.

Parallelization. Microscope time is the most valuable resource. The framework offloads POI identification and region calculation to a separate process, keeping the microscope continuously occupied. The efficiency of this approach depends on the native feature density and the speed of POI identification.

2.3. User workflow

To apply the framework to a new use case, users require two components: a microscope interface for their hardware providing the required methods described above, and a detection function that receives a fast scan image and returns regions that are to be rescanned. Before scanning, it is recommended to evaluate the detection method on recorded image pairs acquired with θ_{fast} and θ_{analysis} to verify that POIs are reliably identified in the fast scan images and to determine suitable fast scan parameters. The framework provides a mode to automatically acquire these paired images, configured through the same configuration file as used for the full scan. Alternatively, if the microscope is not available for preliminary testing, synthetic fast scan images can be generated by degrading existing high-resolution micrographs, for example by downsampling to the target resolution of θ_{fast}. This allows users to develop and evaluate their detection pipeline before committing microscope time.

With these components in place, the user configures the scan and panorama parameters in a YAML configuration file and starts the scan. The interactive setup phase is shown in Fig. 2: the user focuses and adjusts brightness and contrast at three corner positions of the panorama, from which the framework interpolates imaging parameters across the scan area. Three corners are sufficient to define a linear gradient across the panorama, which assumes an approximately planar sample surface with possible tilt. Once the setup is complete, the framework proceeds autonomously through the parallelized scanning workflow.

2.4. Scanning procedure

Efficient imaging requires the microscope to scan arbitrary subregions of the field of view, as stage movements between scans add overhead and introduce misalignment. Each tile is divided into four partitions, and for each partition the process consists of: (1) fast initial scan, (2) POI identification, (3) region proposal, and (4) analysis rescan.

The number of partitions balances two competing effects: more partitions provide more time for the detection process, as regions from earlier partitions can be rescanned while later partitions are still being acquired. However, features that lie on a partition boundary may be split across adjacent partitions and missed by the detection. The probability of this occurring scales with the number of boundary lines and quadratically with the minimum feature size, but inversely quadratically with the tile size in pixels. For the tile and feature sizes used in this work, this probability is negligible. Two partitions

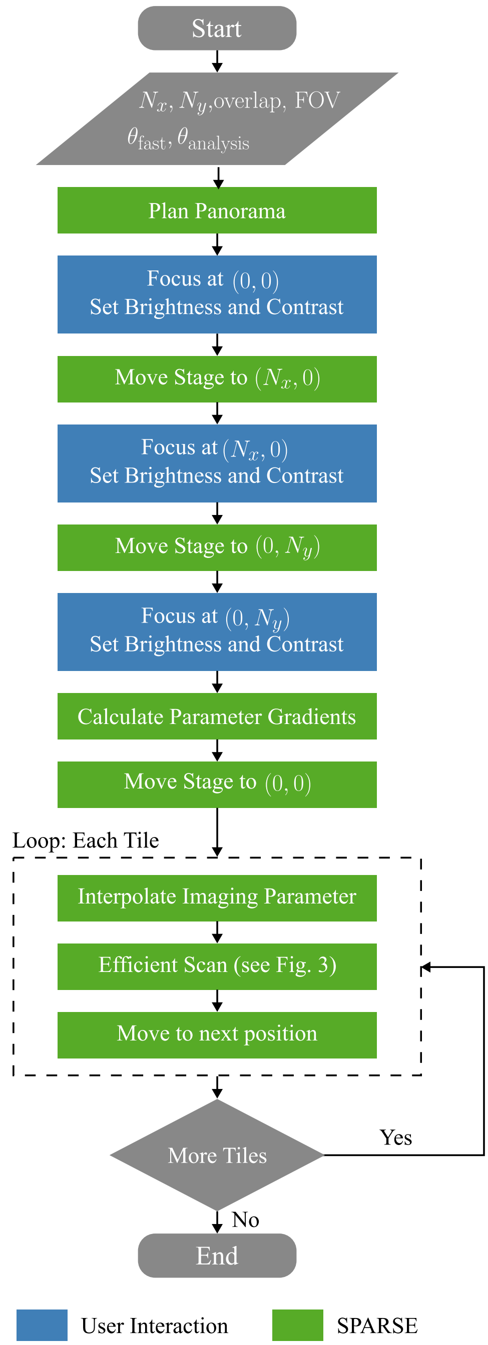

Fig. 2. Flowchart of the scanning procedure. Blue boxes indicate steps requiring user interaction; green boxes indicate automated processes. Users set imaging parameters at three corner positions, from which the framework interpolates parameters across the entire panorama.

would minimize boundary effects but provide limited opportunity for overlapping detection and rescanning with acquisition: the detection must complete and all rescan regions must be acquired before the second partition finishes scanning. With four partitions, the SEM can rescan detected regions from earlier partitions while later partitions are still being acquired and processed. This balances detection time against boundary effects.

These four steps of the scanning procedure are parallelized as shown in Fig. 3. Initially, the microscope scans the first two partitions with θ_{fast}. Once the first partition is acquired, the detection process begins analysing it while the microscope continues with the second partition.

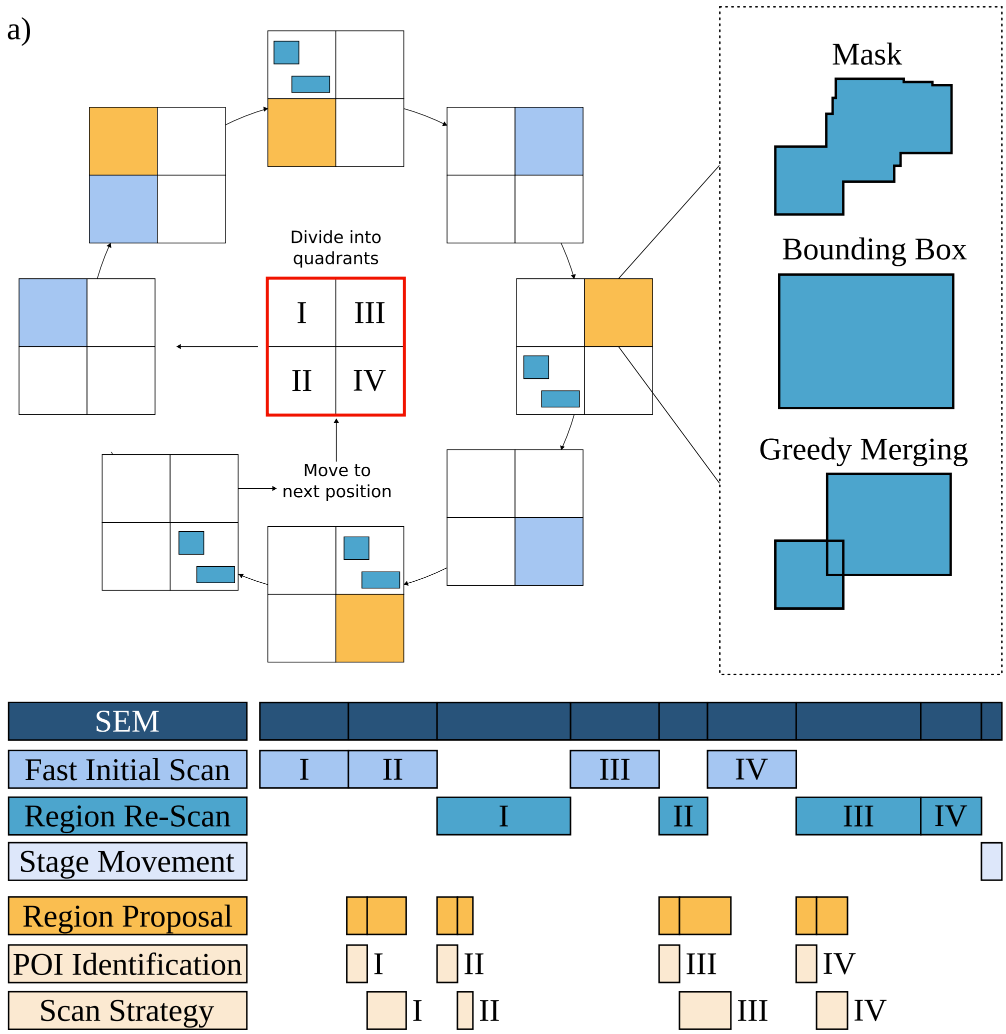

Fig. 3. Parallelized workflow for a single tile divided into four partitions. Blue indicates SEM operations (fast scanning or rescanning); orange indicates concurrent analysis on the workstation. While the microscope performs the fast scan of partition n, region proposal for partition $n-1$ proceeds in parallel. Once the fast scan completes, the microscope rescans detected regions in partition $n-1$ before advancing to the fast scan of partition $n+1$. This pattern ensures analysis time is hidden behind acquisition time.

After the second scan completes, region proposals for the first partition should be ready, allowing the microscope to re-scan the first partition while analysis of the second partition proceeds. This alternating pattern continues until all regions in the final partition have been re-scanned. Then the microscope moves to the next tile. The following subsections describe each step in detail.

2.4.1. Fast initial scan

The first scan is performed using imaging parameters θ_{fast} that result in a faster initial acquisition. The resulting image is used to localize points of interest. The specific imaging parameter choice needs to balance an acceleration in imaging with enough recorded information to still be able to localize points of interest. The fast scan parameters can employ a lower resolution, reduced dwell time, or both.

2.4.2. POI identification

From the fast scan image, points of interest are identified using the selected detection method. This function receives the fast scan image as a NumPy array and returns a list of POI coordinates (see Section 2.5.4 for details). The framework is agnostic to the specific algorithm: threshold-based segmentation [23] or deep learning approaches [4,24] can all be used through the same interface.

2.4.3. Region proposal

Given the identified POIs, rectangular regions of size $w_{\text{region}} \times h_{\text{region}}$ are generated around each point. Since these regions may overlap, scanning each separately would result in redundant acquisition. The framework provides three strategies for combining overlapping regions, and the user selects the appropriate one in the configuration based on their microscope's scanning capabilities. The choice depends on the scanning time per region t_{img}, which is given by:

$$t_{\text{img}} = t_{\text{dwell}} \cdot n_{\text{pixel}} + t_{\text{row}} \cdot n_{\text{rows}} + t_{\text{overhead}} \tag{1}$$

where t_{dwell} is the dwell time per pixel, n_{pixel} is the total number of pixels, t_{row} is the line flyback time, n_{rows} is the number of scan lines, and t_{overhead} is the fixed overhead between scans. The parameters t_{dwell}, t_{row}, and t_{overhead} can either be taken from the microscope specifications or determined empirically by fitting the timing model to measured scan times.

Mask-based merging. If the microscope's scan generator accepts arbitrary masks, all regions are combined into a single binary mask. The user provides the conversion from this mask to the format required by their scan generator as part of the microscope interface. This is the most efficient strategy when available.

Bounding box merging. Overlapping regions are merged into their combined bounding box. This is efficient when t_{overhead} is large relative to the cost of scanning additional pixels. This is the simplest strategy.

Greedy merging. For each pair of overlapping regions, the times for separate scanning t_{separate} versus merged scanning t_{merge} compare as follows:

$$t_{\text{separate}} = t_{\text{dwell}}(n_1 + n_2) + t_{\text{row}}(n_{1,\text{rows}} + n_{2,\text{rows}}) + 2t_{\text{overhead}} \quad (2)$$

$$t_{\text{merge}} = t_{\text{dwell}} \cdot n_{\text{merged}} + t_{\text{row}} \cdot n_{\text{merged,rows}} + t_{\text{overhead}} \quad (3)$$

where n_{merged} and $n_{\text{merged,rows}}$ refer to the merged bounding box dimensions. Regions are merged only if $t_{\text{merge}} < t_{\text{separate}}$. This strategy is most efficient when t_{overhead} is small.

The choice of merging strategy and its parameters therefore constitutes a numerical trade-off between the per-scan overhead and the cost of acquiring additional pixels: the greedy strategy uses the measured timing parameters of Eq. (1), in particular the fixed overhead $t_{overhead}$ and the line time t_{row}, to decide for each pair of regions whether merging reduces the total scan time.

2.4.4. Region rescan

The optimized regions are scanned using θ_{analysis}, the imaging parameters required for subsequent quantitative analysis.

2.5. Implementation details

The framework is implemented in Python 3.11 and runs on a workstation connected to the microscope control PC over the network. The reference implementation communicates with the microscope through the Tescan SharkSEM API. The framework does not depend on any platform-specific libraries apart from the vendor-specific microscope control interface.

The framework is implemented as a Python library using abstract base classes to define the microscope and detection interfaces. Users subclass these and pass instances to the framework at initialization, for example `sem = TescanClara(params)`, where `TescanClara` inherits from the framework's base microscope class. This approach was chosen over alternative architectures for its simplicity. A microservices architecture, where microscope control, detection, and scan planning run as separate services communicating over a network, would allow each component to run on a different machine and in a different language. However, it introduces substantial complexity in deployment, error handling, and inter-service communication that is not justified for a two-process pipeline where low latency between microscope commands is essential. Asynchronous frameworks such as asyncio[1] operate within a single thread and cannot provide true parallelism for CPU-bound detection tasks. Distributed computing frameworks such as Dask [25] would allow scaling detection across multiple machines but add a heavy dependency and scheduling overhead that does not pay off for the current workload. The library approach keeps the barrier to adoption low for researchers familiar with the scientific Python ecosystem, while the abstract base class pattern ensures that new hardware integrations follow a consistent structure.

2.5.1. Parallelization strategy

The parallelized workflow described in Section 2.4 is realized using Python's `multiprocessing` standard library with two processes, the SEM process and the region proposal process, communicating through two shared queues (the concurrency is illustrated in Fig. 3). The SEM process holds the full partition sequence and advances through it autonomously. After acquiring a partition with θ_{fast}, it passes the image to the region proposal queue and checks its own queue for pending rescan regions. If rescan regions are available, they are scanned before advancing to the next partition. If not, the SEM process proceeds to the next partition immediately. The region proposal process receives images, performs POI identification and region proposal, and adds detected regions to the SEM queue for rescanning. Since the SEM process manages its own scanning sequence, it is never idle waiting for the region proposal process to complete. Both partition scans and rescan regions use the same coordinate format, so the SEM process handles them identically. Inter-process communication via queues requires serialization of the transmitted data. The SEM process passes acquired partition images as NumPy arrays to the region proposal queue, while the region proposal process sends scan coordinates back to the SEM queue. Rescanned regions are saved directly to disk by the SEM process and are not passed between processes. The serialization overhead scales with image size but remains negligible relative to acquisition time: even for a 3072×3072 pixel 16-bit image (18 MB), serialization takes on the order of milliseconds, while acquiring such an image at typical dwell times takes seconds to minutes. For applications where this overhead becomes relevant, Python's `multiprocessing.shared_memory`[2] module could be used to avoid copying.

Process-based parallelism is chosen over thread-based concurrency to ensure that the SEM process is never blocked by the region proposal process. In Python, threads share a global interpreter lock (GIL), which prevents concurrent execution of Python code. While many numerical libraries release the GIL during computation, the framework cannot assume this for user-provided detection methods. Separate processes each operate with an independent interpreter, ensuring that a computationally intensive detection step cannot block microscope communication. Recent Python versions (3.13+) introduce experimental free-threaded builds that remove the GIL, which would allow thread-based parallelism for CPU-bound tasks. However, as this feature is not yet the default and library support remains incomplete, the framework uses process-based parallelism to ensure compatibility across standard Python installations. Since the two processes do not share mutable state and communicate exclusively through thread-safe queues, no explicit synchronization mechanisms such as mutex locks are required, and race conditions are avoided by design.

The main process monitors the region proposal process. If it terminates unexpectedly, the main process triggers the shutdown procedure. Queue operations use timeouts to prevent indefinite blocking in case communication is interrupted. During shutdown, whether due to an error or successful completion, the controller writes the index of the last completed tile to disk. Upon restart, the framework reads this index and resumes scanning from the next tile, avoiding the need to repeat already acquired data.

2.5.2. Coordinate systems

Stage movements between tiles use the microscope coordinate system in millimetres. Within a tile, all operations, including detection, region proposal, and rescanning, use pixel coordinates relative to the field of view. Rescanning detected regions does not require stage movement, as the microscope scans arbitrary sub-regions of the field of view by directing the electron beam (see `scan_area` in the microscope interface).

2.5.3. Microscope interface

The framework defines an abstract base class for microscope control. To adapt the framework to a new microscope, users inherit from this class and map the required methods to the corresponding functions in their microscope's control software or API. The implemented class is then passed to the framework at initialization, which handles integration with the scanning controller. The following methods need to be implemented:

[1] https://docs.python.org/3/library/asyncio.html.

[2] https://docs.python.org/3/library/multiprocessing.shared_memory.html.

- `connect()`: Establish a connection to the microscope. The method is expected to raise an exception if the connection fails.
- `disconnect()`: Close the connection to the microscope.
- `prepare()`: Perform preparation steps such as enabling the electron beam and selecting the detector. The specific steps depend on the instrument and are determined by the user. The framework calls this method once before the scanning procedure begins and stores the initial stage position, imaging parameters, and working distance for later restoration.
- `scan_area(coordinates, dwell_time, window_size)`: Scan a specified sub-region of the field of view and return the image as NumPy array. The window size and scan window coordinates are specified in pixels, the dwell time in nanoseconds. This differs from scanning the full FOV and is essential for both the parallelization strategy and targeted rescanning of detected regions (see Section 2.4.4).
- `move_stage(coordinates)`: Move the stage to the specified position. Coordinates are in millimetres in the microscope coordinate system (see Section 2.5.2.
- `reset_microscope()`: Return the stage to its initial position and restore all imaging parameters. If configured, the electron beam is switched off. This method is also called in case of an error to ensure the microscope is left in a defined state.
- `save_metadata(path)`: Save the current stage position, device parameters, and timestamp to the specified path.

Optional methods include:

- `set_working_distance(value)`: Adjust focus. The working distance is specified in millimetres. Required for large-area scans where sample tilt causes focal drift.
- `get_working_distance()`: Retrieve the current working distance in millimetres. Required to determine working distance gradient for large-area scans.
- `set_brightness_contrast(brightness, contrast)`: Adjust detector settings. Values are specified in percent. Required when imaging conditions vary across the scan area.
- `get_brightness_contrast()`: Get brightness and contrast values in percent to estimate brightness and contrast gradients over large-area scans.

2.5.4. Detection interface

The detection method is implemented as a function that receives the fast scan image as a NumPy array and returns a list of POI coordinates in pixels relative to the field of view at the analysis resolution defined by θ_{analysis}. The framework then generates regions around the identified POIs and combines overlapping regions using the merging strategies described in Section 2.4.3. The region dimensions are specified in the configuration file and depend on the requirements of the subsequent analysis, for example the input size of a classification model.

2.5.5. Additional features

Scan resumption. The framework supports resuming interrupted scans using the tile index written during shutdown (see Section 2.5.1). For scans resumed after sample exchange or between microscope sessions, users redefine the panorama corner points in the same order, and the framework recalculates the coordinate mapping.

Parameter interpolation. For large panoramas, imaging parameters (working distance, brightness, contrast) can be set at the corners and linearly interpolated across the scan area to compensate for sample tilt or varying surface conditions.

Error recovery. If any process terminates unexpectedly, the framework calls `reset_microscope()` to return the stage to its initial position and restore all imaging parameters. Currently, the framework handles errors at the process level: if any of the two processes (SEM, or region proposal) throws an unhandled exception, the shutdown procedure is triggered.

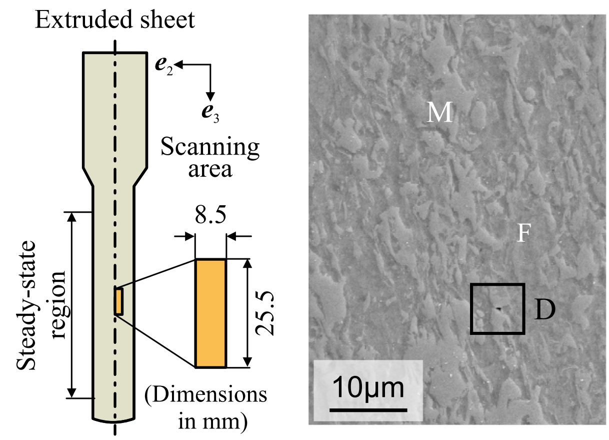

Fig. 4. Schematic representation of the scanned area on the left and a representative micrograph on the right. M marks martensite, F ferrite, and D a damage site.

2.5.6. Data output

Each tile is saved in a separate directory within a root directory specified in the configuration file. A copy of the configuration file is stored in the root directory. Within each tile directory, images are named according to their position and extent in the field of view as `x_y_width_height.tiff`. This naming convention applies to both the fast scan images and the rescanned regions, allowing each image to be located within the tile from its filename alone. The framework saves images as TIFF files with microscope metadata embedded in the file header, including acceleration voltage, probe current, and detector settings. The tile position in the microscope coordinate system is stored in a separate JSON file within the tile directory. Rescanned regions can be linked back to the full panorama by combining their position within the tile with the tile position in the panorama.

2.5.7. Reproducibility

The framework records all information necessary to reproduce a measurement. The configuration file, which specifies the imaging parameters, detection method and its parameters, and panorama coordinates, is saved alongside the output data. To repeat a measurement, users provide the saved configuration file to the framework without further modification. Combined with the scan resumption feature, this allows measurements to be continued across sessions or restarted from scratch with identical parameters.

3. Case study: Damage characterization in dual-phase steel

3.1. Materials and methods

Dual-phase steels are a popular subclass of advanced high strength steels often found in structural applications [1]. As such, their prominence is strongly tied to their mechanical properties, especially energy absorption and component life-time, which are impacted negatively by increased damage prevalence. Dual-phase steels are composed of a ferritic matrix, compounded with martensite islands, and the observed damage phenomena are usually tied to either fracture of martensite or delamination at the interface between martensite and ferrite. In order to validate the proposed framework for efficient damage detection, formed specimens with damage sites distributed across large flat areas are preferred. Therefore, the material is plastically formed by a so-called sheet extrusion process to generate damage sites (process details can be found in [26]). The initial DP800 sheet is positioned between two billet halves made of a bulk material 16MnCrS5 with similar strength. Subsequently, the assembled workpiece is cold extruded using a triple-acting hydraulic press (SMG HZPUJ 260/160). This forming process offers the advantage of homogeneous deformation at large plastic strains. In addition, the extruded sheet exhibits a steady-state region due to its quasi-stationary process characteristic. Moreover, the process is symmetric with respect to the $\boldsymbol{e}_1$-$\boldsymbol{e}_2$ and $\boldsymbol{e}_2$-$\boldsymbol{e}_3$ planes. Therefore, half of the specimen width at any position within the steady-state region can be considered representative for the entire extruded part and is used as the scanning area.

Samples were cut from the extruded sheets by electrical discharge machining (EDM). The cross-sections of interest were then prepared using standard metallographic procedures to obtain a flat, scratch-free surface suitable for SEM observations. For this purpose, the specimens were successively ground with SiC abrasive papers up to a grit size of 4000, followed by polishing on DAC cloths using water-based diamond suspensions with particle sizes of 3 μm and 1 μm. A final chemo-mechanical polishing step with an oxide polishing suspension (OPS, nominal particle size 0.25 μm) was carried out to achieve a mirror-like finish. Subsequently, the polished surfaces were etched in 2 % nital for 6 s, revealing the dual-phase microstructure and providing sufficient contrast to distinguish between martensitic and ferritic regions.

The polished samples were mounted in the SEM (Tescan, Clara) and imaged with an Everhart–Thornley secondary electron detector at a working distance of 10 mm, an accelerating voltage of 20 kV, and a beam current of 3 nA. A schematic representation of the sample and the scanned region, together with an exemplary micrograph are shown in Fig. 4. To enable systematic evaluation of detection performance, a validation dataset was acquired by recording a full panorama using the analysis imaging parameters θ_{analysis} without the automated detection and rescanning workflow. The panorama consisted of 10 × 30 tiles, each covering a field of view of 100 μm × 100 μm, with 15 % overlap between adjacent tiles. Each tile was imaged twice in immediate succession with different parameter sets, minimizing drift between acquisitions. The fast scan parameters θ_{fast} used a dwell time of 32 μs at 768 × 768 pixels, yielding a resolution of 130 nm per pixel. The analysis parameters θ_{analysis} used a dwell time of 32 μs at 3072 × 3072 pixels, yielding a resolution of 33 nm per pixel. After acquiring both images at one position, the stage moved to the next tile. The total imaged area was 2.1675 mm^2.

3.2. Detection pipeline and reference configuration

In dual-phase steels, damage sites manifest as localized clusters of dark pixels in scanning electron micrographs due to the reduced signal intensity at voids and cracks. To detect such features, Kusche et al. [2] proposed a two-stage approach: first, pixels with intensity below a threshold τ are identified as damage candidates, and then spatially coherent groups of candidates are distinguished from noise using density-based clustering (DBSCAN) [27]. DBSCAN assigns each candidate pixel to a cluster if at least $n_{\min}$ other candidate pixels lie within a radius of ε pixels; candidates that do not meet this density criterion are discarded as noise. Each resulting cluster is considered a potential damage site.

Since the detected damage population depends on the choice of detection parameters, we evaluated framework performance across a range of configurations to ensure that the results are not specific to a single parameter set. The parameter ranges were chosen to cover a broad space of possible detection outcomes, varying both the minimum spatial extent and the required local density of damage clusters. The threshold τ defines the maximum grayscale intensity below which a pixel is classified as a damage candidate. Values of $\tau \in [45, 50, 55, 60]$ were selected based on visual inspection of the analysis-scan micrographs, spanning the range from conservative detection of only the darkest voids to more permissive inclusion of lighter damage features. For DBSCAN clustering, the clustering radius ε and minimum neighbours $n_{\min}$ are coupled parameters: therefore we scaled $n_{\min}$ with ε, testing $\varepsilon = 1$ with $n_{\min} \in [4, 5]$, $\varepsilon = 2$ with $n_{\min} \in [10, 13]$, and $\varepsilon = 3$ with $n_{\min} \in [15, 20, 25]$.

3.3. Matching methodology

The following evaluation assesses the detection performance of one specific detection pipeline used with the framework: a fast scan at reduced resolution with bicubic and Lanczos [28] interpolation to the analysis resolution, followed by DBSCAN-based detection. The detection performance depends on the choice of θ_{fast}, upsampling method, and detection algorithm, and is independent of the framework itself.

To evaluate detection performance, we compared damage sites detected in the θ_{fast} images against reference sites identified in the θ_{analysis} images. To establish this correspondence, matching was performed using pixel overlap rather than centroid distance. Centroid-based matching pairs a reference and a detected cluster when the distance between their centroids falls below a fixed threshold. This approach accounts for the fact that processing can alter cluster morphology: a single fragmented cluster in the high-resolution image may appear as a connected region after interpolation, or vice versa. With centroid-based matching, if two nearby clusters in the reference image are detected as a single merged cluster, only one reference cluster would be matched and the other would be counted as a false negative, even though the proposed rescan region covers both. Overlap-based matching instead counts a reference cluster as detected whenever any detected cluster shares at least one pixel with it, correctly identifying both clusters in this case.

The matching procedure was as follows. First, reference damage sites were identified in the θ_{analysis} images using thresholding followed by DBSCAN clustering. For each reference cluster, we recorded the set of pixel coordinates belonging to that cluster. Second, detection candidates were identified in the θ_{fast} images using the same algorithmic pipeline with varying threshold and clustering parameters. Third, for each reference cluster, we checked whether any detection cluster shared at least one pixel with the reference cluster. A reference site was marked as detected if such overlap existed. This permissive criterion is appropriate because the evaluation targets detection, not segmentation: any overlap means a rescan region of size $w_{\text{region}} \times h_{\text{region}}$ is placed around the detected cluster, which is sufficient to acquire an analysis quality image of the reference site. Metrics such as the Intersection over Union, $\text{IoU} = |A \cap B|/|A \cup B|$, or Dice [29], $\text{Dice} = 2|A \cap B|/(|A| + |B|)$, where A and B denote the detected and reference cluster areas, would additionally penalize differences in cluster shape and position between the fast and analysis images, which are not relevant for the rescanning task.

The paired images were acquired sequentially at each tile position, with the θ_{analysis} image recorded immediately after the θ_{fast} image. Despite this, minor stage drift between acquisitions can introduce spatial offsets between corresponding features. To account for this, we dilated each reference cluster by 10 pixels prior to overlap checking. This dilation compensates for alignment uncertainty between the two separately acquired images and does not affect the overlap criterion itself: a detection must still overlap with the dilated reference cluster to count as a match. The dilation is deliberately permissive: rather than estimating the drift, we chose a value large enough that even if it produced a match to a neighbouring feature, the resulting rescan window would still place that feature within its central region. This reflects the purpose of the evaluation, which is to determine whether a feature would be rescanned with sufficient surrounding context to be analysed, not whether a specific detected cluster coincides with a specific reference cluster. Since the rescan patches are much larger than the dilation, a feature counted as detected is captured well within the rescan window. The matching procedure is illustrated for representative damage sites at the 90 % operating point in Fig. B.2. Detection rate was calculated as the fraction of reference sites with at least one overlapping detection:

$$R_{\text{detection}} = \frac{N_{\text{overlap}}}{N_{\text{reference}}} \tag{4}$$

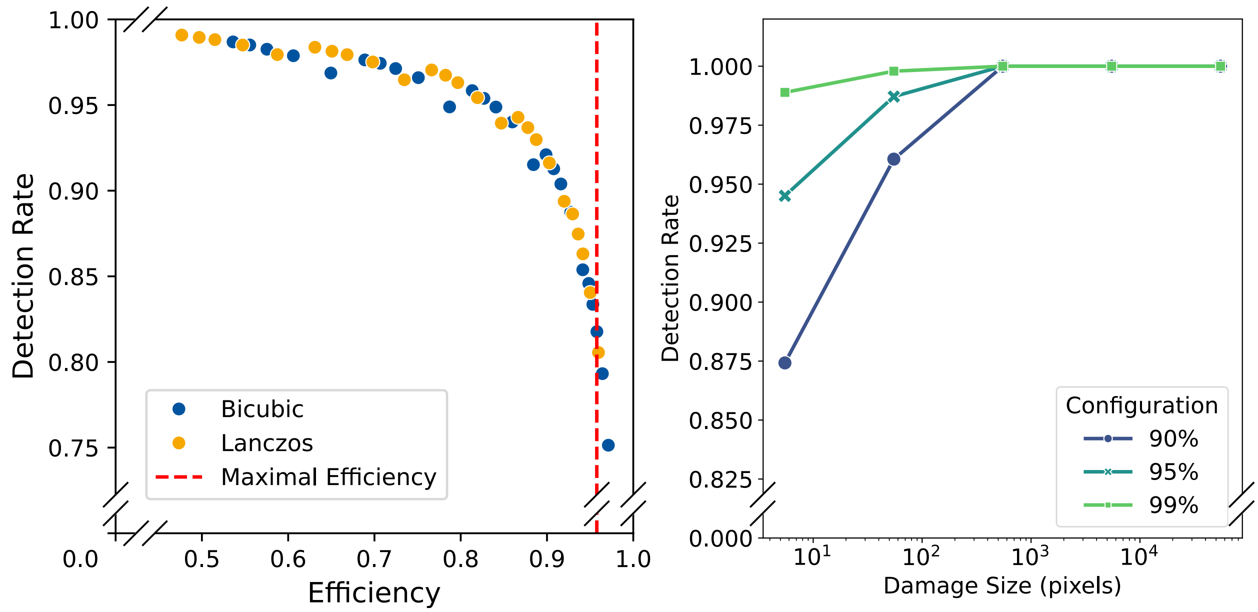

Fig. 5. Detection performance for $\tau = 55$, $\epsilon = 1$, $n_{min} = 5$. (a) Pareto front showing the trade-off between detection rate and efficiency. The dashed line indicates the theoretical maximum efficiency assuming perfect detection of all ground truth clusters. (b) Detection rate as a function of cluster size for three operating points corresponding to 99%, 95%, and 90% overall detection rate. Lines connecting data points added to guide the eye.

where $N_{overlap}$ is the number of reference clusters overlapping with at least one detection cluster, and $N_{reference}$ is the total number of reference clusters.

Efficiency was defined as the fraction of the total scan area that does not require rescanning:

$$E = 1 - \frac{A_{rescan}}{A_{total}} \quad (5)$$

where A_{rescan} is the area covered by rescan patches around detected clusters and A_{total} is the total scanned area. An efficiency of 100 % would indicate that no regions were rescanned, meaning no POIs were detected. In practice, efficiency is bounded from above by the theoretical maximum corresponding to perfect detection of all ground truth POIs with no false positives. This area-based metric is independent of microscope parameters such as dwell time, allowing comparison across different imaging conditions.

3.4. Accuracy-efficiency trade-off

Fig. 5a shows the trade-off between detection rate and efficiency for a representative parameter configuration ($\tau = 55$, $\epsilon = 1$, $n_{min} = 5$). These parameters were chosen for the detection of damage in DP800 under the specific imaging conditions used in this study. The threshold $\tau = 55$ on the 8-bit grayscale range captures the dark contrast of damage sites while excluding the ferritic matrix and depends on the imaging parameters such as brightness, contrast, and detector settings. The clustering parameters $\epsilon = 1$ and $n_{min} = 5$ ensure sensitivity to small and elongated crack-like features. The Pareto front follows the expected trend: at the highest detection rate of 99 %, efficiency is approximately 48 %, meaning roughly half of the scanned area would require rescanning. Relaxing detection requirements improves efficiency substantially—at 95 % detection rate, efficiency increases to 87 %, and at 90 % detection rate, efficiency reaches 92 %. The theoretical maximum efficiency of 96 %, corresponding to perfect detection of all ground truth clusters with no false positives, provides an upper bound for this configuration.

These efficiency gains translate directly into reduced acquisition time. The total acquisition time of the detection pipeline presented here, t_{SPARSE} relative to that required for conventional full-area high-resolution imaging, $t_{\theta_{analysis}}$ can be estimated as

$$\frac{t_{SPARSE}}{t_{\theta_{analysis}}} = \frac{1}{16} + \frac{A_{rescan}}{A_{total}} = \frac{17}{16} - E \quad (6)$$

where the factor 1/16 reflects the reduced pixel count of the fast scan relative to the high-resolution scan, and the second term accounts for the rescanned area. For the operating points discussed above, this corresponds to approximately 58 % of the conventional acquisition time at 99 % detection rate, 19 % at 95 %, and 14 % at 90 %. In addition to the detection rate, we evaluated the precision, the number of false-positive detections, and the number of rescan regions at the three operating points. Computed at the level of detected clusters, consistent with the detection rate, the precision was 0.04, 0.25, and 0.41 at the 99 %, 95 %, and 90 % operating points, with corresponding false-positive counts of 281 440, 29 279, and 13 444. Even conservative operating points with near-complete detection thus yield substantial time savings, while accepting moderate reductions in detection rate can reduce acquisition time by an order of magnitude. For example, acquiring 1 mm^2 with $\theta_{analysis}$ at a dwell time of 32 µs and a resolution of 3072 pixel per 100 µm requires approximately 12 h. At the conservative operating point of 95 % detection rate, this reduces to approximately 2.3 h, and at 90 % detection rate to approximately 1.7 h. These figures are based on the ratio of scanned pixels and do not include instrument overhead.

Fig. 5b shows the detection rate as a function of cluster size for three operating points along the Pareto front (99 %, 95 %, and 90 % overall detection rate). Detection rate generally increases with cluster size, as larger damage sites produce stronger signals in the fast scan images. Detection rate is lowest for the smallest clusters and increases with cluster size, as larger damage sites produce stronger signals in the fast scan images.

3.5. Sensitivity to detection parameters

The full evaluation across all parameter configurations is shown in Fig. A.1. The results demonstrate consistent behaviour across the tested parameter space: all configurations produce similar Pareto front shapes, with the primary difference being the range of achievable detection rates and efficiencies. Importantly, each parameter configuration defines its own reference population: the ground truth clusters are identified in the θ_{analysis} images using the same parameters, so the detection rate reflects how well the fast scan reproduces the results of the respective full-resolution analysis. Configurations with stricter clustering requirements (larger ε and $n_{\min}$) target larger and denser clusters, which produce stronger signals in the fast scan images and are therefore more reliably detected. This shifts the Pareto front towards higher detection rates and higher efficiencies, but over a smaller population of damage sites. Conversely, configurations with more permissive clustering parameters detect smaller and sparser features, resulting in lower baseline detection rates and reduced efficiency due to the larger number of rescan candidates. These results indicate that the framework performance is robust across reasonable parameter choices, with the specific configuration primarily determining which subset of the damage population is targeted.

4. Discussion

4.1. Comparison to prior work

Recent work by Clusiau et al. demonstrated workflow automation for SEM acquisitions with feature tracking for nanoparticle analysis [17]. Our framework shares the goal of reducing acquisition time through selective high-resolution imaging but differs in two main aspects. First, our implementation parallelizes scanning, detection, and rescanning, decoupling computation time from acquisition time. In a sequential workflow, any time spent on detection directly adds to the total acquisition time, which limits users to simple and fast detection methods. With our parallel approach, more demanding algorithms can run alongside scanning without penalty, as long as processing finishes before the next partition is ready. If this constraint becomes limiting, the modular architecture allows computation to be distributed across multiple processing instances, ensuring that acquisition speed remains the bottleneck rather than detection. Second, we systematically quantify the accuracy-efficiency trade-off across detection parameters, allowing users to choose an operating point that balances detection completeness against acquisition time.

4.2. Detection performance in dual-phase steel

For the selected parameter configuration ($\tau = 55$, $\varepsilon = 1$, $n_{\min} = 5$), nearly all damage site candidates could be identified, achieving a detection rate of 99 %. However, this high detection rate comes at a significant cost: the rescan mask covers approximately half of the total area. For microscopes without the capability to scan arbitrary masked regions, this would require rescanning large bounding boxes, potentially negating the time savings. Fig. 6 illustrates this trade-off: at 99 % detection rate, the surface is densely covered with rescan patches, while at 95 % detection rate, only isolated regions require rescanning.

Furthermore, detected POIs are not necessarily damage sites, as some correspond to imaging artifacts such as shadows or contrast variations. This is a known challenge in damage characterization of dual-phase steels: Medghalchi et al. [6] addressed this by including a shadow class alongside the damage mechanism classes in the classification task. In principle, the framework supports a workflow where detected candidates are first rescanned at analysis quality and then classified using a trained model that distinguishes damage mechanisms from artifacts. However, implementing and evaluating such a classification pipeline is beyond the scope of this work.

The chosen DBSCAN configuration ($\varepsilon = 1$, $n_{\min} = 5$) targets small, dense clusters and represents a challenging case for interpolation-based detection. We chose this configuration for the main analysis because the small, difficult-to-detect clusters it targets act as a stress test of the method: larger damage sites are expected to be detected with the same or better performance. As shown in Figure A1, this expectation is confirmed: other physically reasonable configurations achieve substantially better performance. For example, with $\tau = 55$, $\varepsilon = 2$, and $n_{\min} = 13$, the detection rate reaches 100 % with efficiency above 90 %, corresponding to an estimated acquisition time of approximately 16 % of the conventional full-area high-resolution approach.

The practical value of simple interpolation-based detection depends on the microscope's scanning capabilities. When mask-based scanning is available, the framework achieves substantial time savings even with moderate false positive rates, as only the masked pixels are rescanned. In this case, simple interpolation methods suffice for identifying most ground truth damage sites while still accelerating acquisition. However, when only bounding box or greedy merging strategies are available, false positives have a larger impact on efficiency: each false positive enlarges the rescan region, and at high detection rates the bounding boxes may cover most of the tile area (Fig. 6). The precision of the detection pipeline is low at high detection rates and improves as the operating point is relaxed. This reflects the permissive detection parameters required to capture nearly all features: at the most permissive setting almost every pixel below the threshold forms a candidate cluster, the majority of which are noise rather than damage. The large number of false-positive clusters does not translate proportionally into rescanned area, since the spurious clusters are predominantly small and overlapping detections are merged before rescanning. Because acquisition time is governed by the rescanned area rather than the number of detected clusters, the low cluster-level precision has a limited effect on the achievable efficiency. In these cases, detection methods that reduce false positives while maintaining high detection rates become essential. Learned approaches may provide the necessary improvement in precision to make the framework practical on hardware with more limited scanning flexibility. For example, in previous work [30], we demonstrated that a reference-based super-resolution network trained on dual-phase steel micrographs outperforms standard interpolation in preserving microstructural features relevant to damage detection, such as phase boundaries and voids. Due to the modular detection interface, integrating such a method requires only replacing the detection function, without modifications to the scanning workflow or parallelization strategy. This could improve damage identification in the enhanced fast-scan images, reducing false positives while maintaining high detection rates. More generally, detection methods with higher precision, such as convolutional neural networks, may further reduce acquisition time. Because detection runs in parallel with acquisition, the detection step does not extend acquisition time as long as processing completes before the corresponding rescan decision is required. Improved precision then reduces the number of false positives and therefore the rescanned area, which is particularly relevant on instruments with a large per-scan overhead. The achievable fast-scan resolution is bounded by the requirement that features remain detectable; for the small damage sites studied here this bound is set by the feature size. For larger or higher-contrast features, the fast scan can be acquired at lower quality, either at lower resolution, where the overview contribution to the acquisition time scales as $1/r^2$ with the linear resolution ratio r, or at shorter dwell time, where it scales inversely with the dwell-time reduction, or both. This applies to detection methods able to identify features in such degraded fast-scan images, whether by enhancing the fast scan before detection, for example using the reference-based super-resolution approach described above, or by employing a detector trained to localize features directly at reduced quality.

4.3. Implications for statistical analysis of rare events

Statistically meaningful analysis of rare microstructural features requires scanning areas far larger than typically feasible with conventional high-resolution imaging. In dual-phase steels, previous work has established methods for large-area damage characterization and classification [2,6], three-dimensional damage analysis [31], and resolution

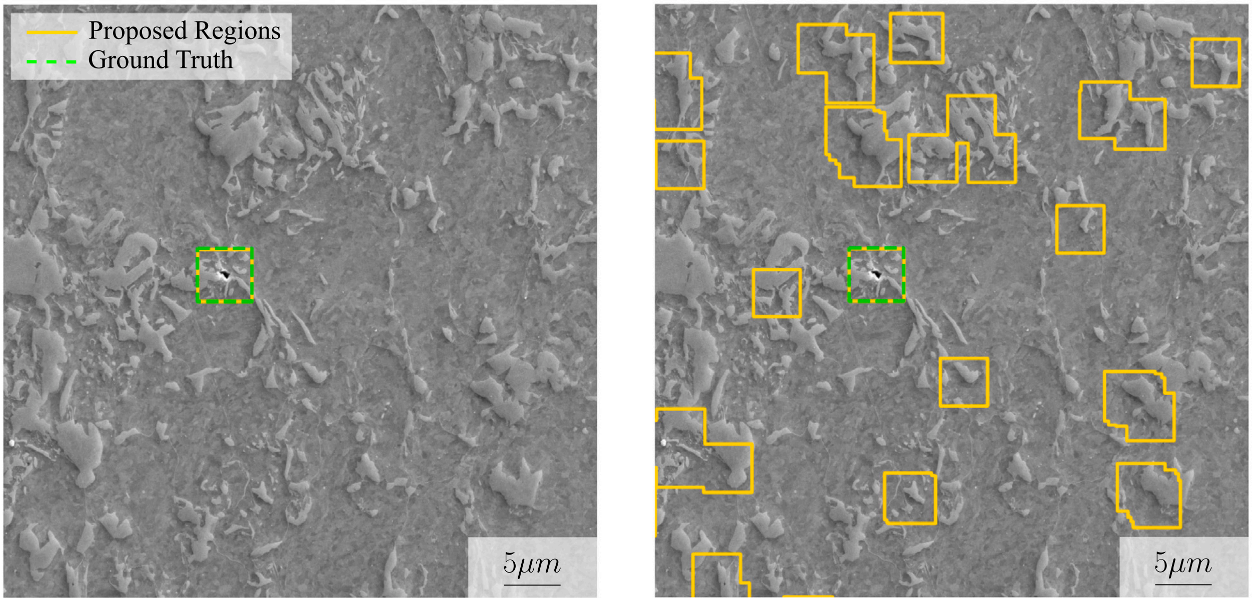

Fig. 6. Comparison of rescan regions for two operating points on the Pareto front. Outlines of proposed rescan regions are shown in yellow, ground truth damage sites in green. Left: at 95 % detection rate (87 % efficiency), only isolated regions require rescanning. Right: at 99 % detection rate (48 % efficiency), the rescan regions cover approximately half of the shown area.

enhancement [30]. However, the acquisition time required for high-resolution panoramic imaging remains a key bottleneck that limits the areas that can be practically examined.

The framework presented here directly addresses this bottleneck by substantially reducing the acquisition time per unit area. Crucially, these time savings can be reinvested: at 10 % of the conventional acquisition time approximately ten times the area can be scanned in the same time budget. Alternatively, the same area can be scanned at multiple strain increments during an in-situ experiment, rather than relying on a single post-mortem analysis. The importance of large-area analysis has been demonstrated by Medghalchi et al. [32], who showed that martensite phase fraction measurements in the same DP800 steel exhibit large scatter even for areas exceeding 1 mm^2, due to long-range microstructural heterogeneity such as martensite banding. Their analysis revealed that window sizes commonly used in the literature are insufficient for representative measurements, and that the associated uncertainty often exceeds the precision at which phase fractions are reported. While their work focused on phase fraction, the same argument applies to damage statistics: reliable characterization of spatially heterogeneous features requires surveying areas substantially larger than current practice, which is only feasible if the acquisition time per unit area can be significantly reduced.

For the dual-phase DP800 steel studied here, the choice of operating point can be guided by the downstream analysis requirements. At 99 % detection rate, nearly all damage candidates are captured, but acquisition still requires approximately 58 % of the conventional imaging time. This operating point is appropriate when completeness is critical, for example when characterizing the full damage size distribution including the smallest feature. At 95 % detection rate, acquisition time drops to approximately 19 % of the conventional approach. The missed detections are predominantly the smallest clusters (Fig. 5b), which contribute least to the overall damage area and are most likely to include false positives. When applying a minimum cluster size of 20 pixel for damage classification in DP800, only three candidates are missed at this operating point, none of which correspond to actual damage sites upon manual inspection. For analyses focused on damage morphology, spatial distribution, or correlation with microstructural features, this represents a practical sweet spot: the statistical power gained from scanning larger areas more than compensates for the loss of a small number of marginal detections. At 90 % detection rate, acquisition time is reduced to approximately 14 %, enabling rapid surveys of very large areas at the cost of missing some smaller damage sites.

The acceptability of missing small damage candidates depends on the analysis objective. The absolute lower bound on detectable damage size is set by the analysis resolution, which is chosen to resolve the mechanisms underlying damage nucleation. Beyond this limit, the physically relevant size range is constrained by the application: unless an investigation of nucleation mechanisms is the specific objective, those damage sites may considered most important that are closer to becoming critical by growth and coalescence and therefore affect the macroscopic failure behaviour of the material, while very small sites must first reach a certain size to become mechanically relevant. As long as the chosen detection threshold is applied consistently across experiments, quantitative comparisons remain valid regardless of where exactly this threshold is set.

The operating points reported here are specific to the detection of damage in DP800 with the thresholding and clustering pipeline. For other materials, feature types, and detection methods, the achievable acceleration differs: it is governed by the fraction of area that must be rescanned, which depends on the feature density, and by the precision of the detection method, which determines how many false positives enlarge that fraction. This is particularly relevant for features that differ from damage voids in contrast or morphology: the thresholding and clustering pipeline used here exploits the strong dark contrast and compact shape of voids and cracks, whereas features such as fine precipitates, non-metallic inclusions, or phase boundaries present different contrast and geometry and would generally require a detection method suited to them rather than intensity thresholding, together with fast-scan parameters that preserve the relevant contrast. The framework and its parallelization remain unchanged; only the detection function and, with it, the position of the accuracy-efficiency trade-off shift.

The acquisition time of the framework is determined by a small set of parameters, which makes the expected performance predictable for a given instrument and application without requiring instrument-specific calibration of the framework itself. The total time comprises the fast overview of the full area and the subsequent rescanning of detected regions. The overview time is fixed by the fast-scan resolution and the per-pixel dwell time, together with the line flyback and per-scan overhead of the instrument. The rescan time depends on the feature density, the false-positive rate of the detection method, and the merging strategy used to combine detected regions into scans. The pixel-ratio estimate of Eq. (6) is the limiting case of this relation when the per-scan overhead is negligible; for instruments with appreciable overhead, the realized time exceeds this estimate. The achievable acceleration for a given application can therefore be estimated from the instrument timing parameters and the feature density and detection characteristics. Of these factors, the framework directly influences only the detection method and thus the false-positive rate; the feature density is set by the material, and the per-scan timing by the instrument and its vendor control software. The savings reported here therefore arise from acquiring fewer high-resolution pixels, on top of the per-scan performance provided by the instrument, rather than from any change to the underlying scan hardware. To illustrate this dependence, Appendix C tabulates the estimated time ratio across a representative range of flyback and per-scan overhead values for both the greedy and masked merging strategies.

The same considerations apply to the volume of stored data, which scales with the number of stored pixels. When the fast scan uses a lower resolution, the overview occupies a fraction $1/r^2$ of a full-resolution image for a linear resolution ratio r, and the total stored data is $1/r^2$ plus the rescanned fraction of the conventional volume, the same expression as the idealized time ratio in Eq. (6). For the four-fold resolution ratio used here, this corresponds to approximately one sixteenth of the conventional data volume plus the rescanned regions. When the fast scan instead uses the full resolution with a reduced dwell time, the overview is the same size as a conventional full-resolution image, so the total stored data exceeds the conventional volume by the rescanned fraction. Reducing the dwell time therefore accelerates acquisition without reducing the stored data volume, whereas reducing the resolution reduces both.

Beyond the dual-phase steel studied here, the framework's applicability extends across imaging modalities: The framework requires a programmable stage and a programmable acquisition region, and its benefit depends on how the imaging modality acquires data. Where acquisition time scales with the imaged area, either in point-scanning modalities such as confocal or scanning probe microscopy, or in large-area tiled acquisition where many fields of view are recorded sequentially with stage movement, the same two-stage strategy of a fast overview followed by selective high-resolution acquisition of detected regions can be applied, with the detection function adapted to the relevant feature contrast. In modalities that capture an entire field of view in a single exposure, by contrast, the time saving arises at the level of selecting which fields to revisit at high resolution rather than within an individual image.

4.4. Practical recommendations

Based on the evaluation presented here, several observations may inform users applying the framework to new materials and feature types:

1. *Simple methods as a baseline.* Interpolation-based detection requires no training data and has provided a strong baseline in this study. The achievable efficiency depends on the material, microstructure, and feature size, and should be evaluated before investing in more complex methods.
2. *Paired validation data.* Acquiring a small paired dataset with both θ_{fast} and θ_{analysis} images enables quantification of the detection performance for a given combination of material, imaging parameters, and detection method. This allows informed selection of operating points on the Pareto front.
3. *Feature size relative to resolution.* In this study, features near the resolution limit of the fast scan showed reduced detection rates. Users working with similarly small features may need to adjust the resolution ratio or detection method accordingly.
4. *Detection rate over efficiency.* In the context of this framework, detecting regions that do not contain features of interest only cost additional acquisition time, while missing actual features leads to incomplete data. Optimizing for high detection rate at the cost of reduced efficiency is therefore generally preferable.

5. Conclusion

We presented an open-source Python framework for efficient large-area SEM characterization of rare microstructural features. This SPARSE framework decouples detection from acquisition through a generic microscope interface that separates microscope-specific control from the scanning logic. This design is intended to allow deployment across different microscope platforms without modifications to the core framework. Scanning, detection, and rescanning are parallelized using separate processes, ensuring that computation time is hidden behind acquisition time. Users can select their operating point on the accuracy-efficiency trade-off based on application requirements.

Validation on damage detection in dual-phase DP800 steel demonstrated that simple interpolation-based detection achieves substantial time savings without requiring training data. Detection rates above 95 % were maintained while reducing acquisition time to approximately 19 % of the conventional full-area high-resolution approach based on the ratio of scanned pixels. For the specific damage characteristics and density observed in this DP800 steel at the given degree of deformation, and the minimum cluster size established by [2] for damage classification, no actual damage sites are missed at this operating point. For stricter clustering parameters, which target larger and more prominent damage sites, all ground truth candidates were detected while reducing acquisition time to approximately 16 %. The trade-off behaviour was robust across a wide range of detection parameters, with stricter configurations identifying smaller features at the cost of reduced efficiency. These time savings can be reinvested into scanning larger areas, directly improving the statistical power of rare-event analyses.

The detection method presented here serves as a baseline to validate the framework. The modular detection interface allows it to be substituted with different approaches, such as resolution enhancement [30], localization directly in the initial images, or denoising, without modifications to the scanning workflow. The framework also provides practical features for extended imaging campaigns, such as scan resumption across sessions, parameter interpolation, and failure recovery.

The complete implementation is available at https://doi.org/10.5281/zenodo.20321819.

CRediT authorship contribution statement

Tom Reclik: Writing – original draft, Software, Methodology, Investigation, Formal analysis. **Jan Gerlach:** Writing – review & editing, Methodology, Investigation. **Maximilian A. Wollenweber:** Writing – review & editing, Methodology, Investigation. **Yannis P. Korkolis:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization. **Sandra Korte-Kerzel:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization. **Ulrich Kerzel:** Writing – review & editing, Supervision, Funding acquisition, Conceptualization.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this work the authors used the LLMs RWTHGPT and Claude (claude.ai) in order to improve language and readability. After using this tool/service, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors express their gratitude to the Deutsche Forschungsgemeinschaft (DFG) for financial support in context of the Collaborative Research Centre CRC/Transregio 188/2 “Damage Controlled Forming Processes”, project no. 278868966, and the Collaborative Research Centre CRC 1394 “Structural and chemical atomic complexity– from defect phase diagrams to material properties”, project no. 409476157. This project has also received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 Research and Innovation Programme (Grant Agreement No. 101168203 “TailorPlast”). The authors thank TESCAN GROUP a.s. for their support and collaboration.

Appendix A

See Fig. A.1.

Appendix B

See Fig. B.2.

Appendix C

The acquisition time of the framework relative to conventional imaging depends on the instrument timing parameters, in particular the line flyback time t_{row} and the per-scan overhead $t_{overhead}$ of Eq. (1). As these vary across instruments and scan-generator implementations, we do not report a single measured value but instead tabulate the estimated time ratio across a representative range, so that the expected acceleration can be read off for a given instrument. The raw detections per tile are taken from the DP800 evaluation at each operating point. For masked scanning the resulting geometry is timing-independent. For greedy merging, the region merging is re-solved at each (t_{row}, $t_{overhead}$) setting, since the greedy merge criterion itself depends on these parameters; the number of merged regions, rescanned pixels, and rows therefore vary across the greedy table.

Tables C.1 and C.2 report the estimated time ratio for the greedy merging and masked scanning strategies. The two strategies differ in how the per-scan overhead enters. Under greedy merging, each merged region is acquired as a separate scan, so the overhead is incurred once per region; at operating points that produce many regions, the realized time grows substantially as the overhead increases. In unfavourable configurations, where a large fraction of the area is rescanned and the per-scan overhead is high, the fast overview adds acquisition time rather than saving it, and the ratio can exceed one (Table C.1). In such cases the overview scan does not pay off, and conventional acquisition is preferable. Under masked scanning, all detected regions are acquired in a single masked scan, so the overhead is incurred only once and the time ratio is nearly insensitive to it.

The masked estimate assumes that the scan generator rasters only the rows containing masked pixels, incurring the per-scan overhead a single time. Depending on the implementation, a scan generator may instead raster the bounding box of the masked region or handle the mask differently, so the masked figures represent a best-case estimate of the achievable saving.

Table C.1
Estimated acquisition time of SPARSE relative to conventional full-area high-resolution imaging, $t_{SPARSE}/t_{\theta_{analysis}}$, for the greedy merging strategy. Each subtable corresponds to one operating point and tabulates the ratio as a function of the line flyback time t_{row} and the per-scan overhead $t_{overhead}$. Values below 1 indicate a time saving. The raw detections per tile are taken from the DP800 evaluation at each operating point; the region merging is re-solved for each combination of timing parameters, since the greedy merge criterion depends on them, so the number of merged regions, rescanned pixels, and rows vary across the table. The timing parameters are swept across a range representative of different instruments and scan-generator implementations.

(a) 99 % detection rate

$t_{overhead}$	t_{row}			
	0.02 ms	0.08 ms	0.4 ms	1.5 ms
1 ms	0.740	0.745	0.773	0.857
10 ms	0.744	0.749	0.777	0.861
100 ms	0.785	0.790	0.815	0.891
1 s	1.01	1.01	1.03	1.07

(b) 95 % detection rate

$t_{overhead}$	t_{row}			
	0.02 ms	0.08 ms	0.4 ms	1.5 ms
1 ms	0.210	0.212	0.222	0.253
10 ms	0.212	0.214	0.223	0.255
100 ms	0.230	0.231	0.241	0.271
1 s	0.362	0.364	0.371	0.395

(c) 90 % detection rate

$t_{overhead}$	t_{row}			
	0.02 ms	0.08 ms	0.4 ms	1.5 ms
1 ms	0.149	0.150	0.156	0.176
10 ms	0.150	0.151	0.157	0.177
100 ms	0.161	0.162	0.167	0.187
1 s	0.250	0.251	0.256	0.272

Table C.2
Estimated acquisition time of SPARSE relative to conventional full-area high-resolution imaging, $t_{SPARSE}/t_{\theta_{analysis}}$, for the **masked scanning** strategy, in which all detected regions are rescanned as a single masked acquisition. Each subtable corresponds to one operating point and tabulates the ratio as a function of the line flyback time t_{row} and the per-scan overhead $t_{overhead}$. The per-scan overhead is incurred once, as the masked region is acquired in a single scan. This is a best-case estimate; see the accompanying text. Values below 1 indicate a time saving.

(a) 99 % detection rate

$t_{overhead}$	t_{row}			
	0.02 ms	0.08 ms	0.4 ms	1.5 ms
1 ms	0.586	0.586	0.589	0.596
10 ms	0.586	0.586	0.589	0.596
100 ms	0.586	0.587	0.589	0.596
1 s	0.591	0.591	0.593	0.600

(b) 95 % detection rate

$t_{overhead}$	t_{row}			
	0.02 ms	0.08 ms	0.4 ms	1.5 ms
1 ms	0.196	0.197	0.200	0.209
10 ms	0.196	0.197	0.200	0.209
100 ms	0.197	0.197	0.200	0.209
1 s	0.202	0.203	0.205	0.215

(c) 90 % detection rate

$t_{overhead}$	t_{row}			
	0.02 ms	0.08 ms	0.4 ms	1.5 ms
1 ms	0.143	0.143	0.146	0.154
10 ms	0.143	0.143	0.146	0.154
100 ms	0.143	0.144	0.146	0.155
1 s	0.149	0.149	0.152	0.160

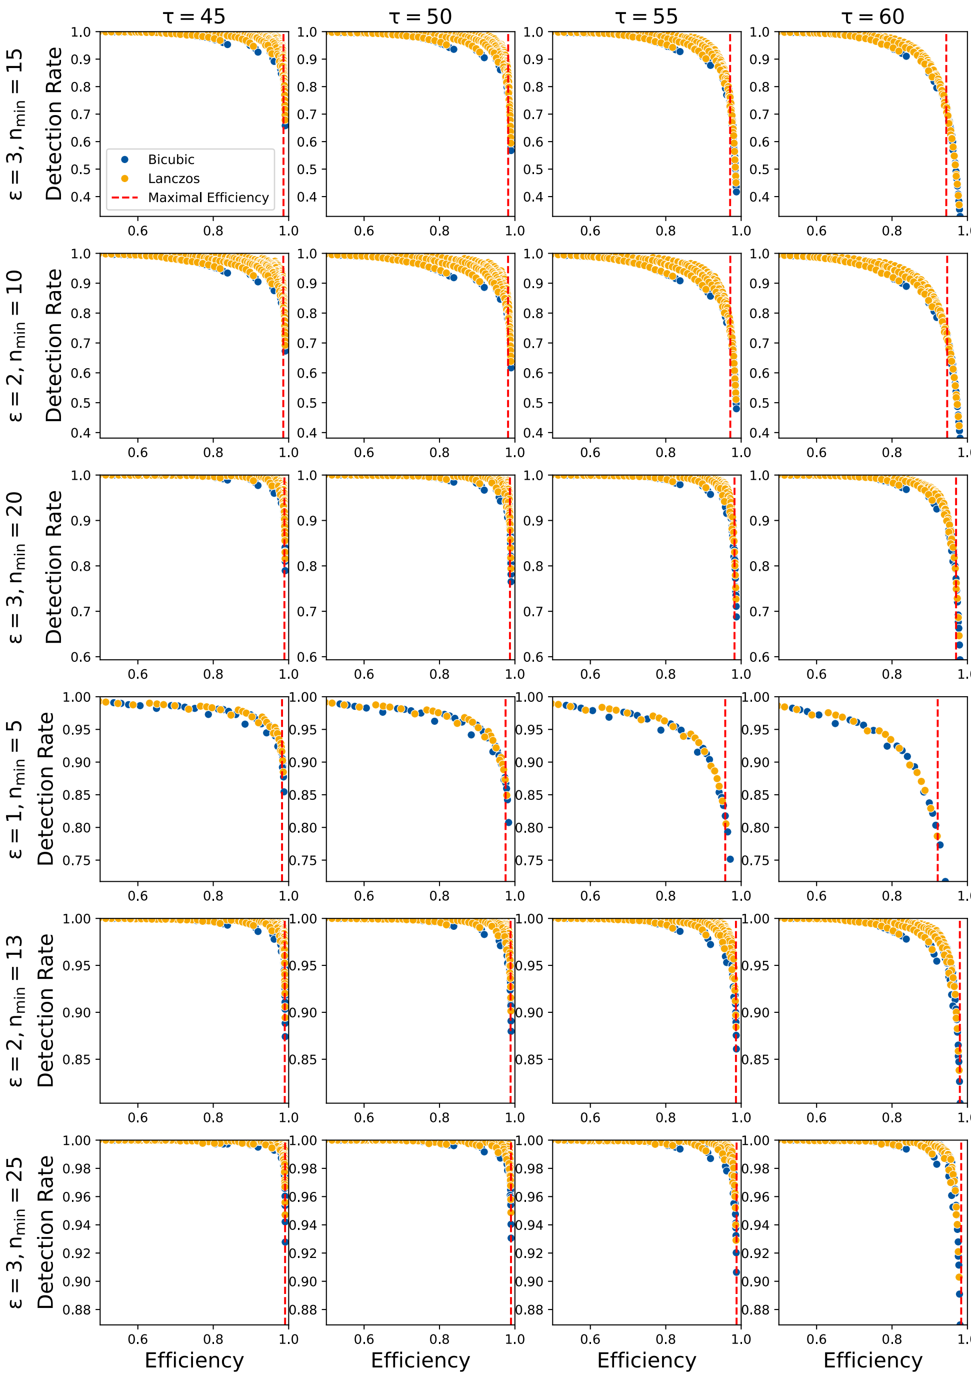

Fig. A.1. Detection rate versus efficiency for all tested parameter configurations. Columns correspond to intensity thresholds $\tau \in [45, 50, 55, 60]$, rows to DBSCAN configurations ordered by increasing cluster density requirements (top: $\varepsilon = 3$, $n_{\min} = 15$; bottom: $\varepsilon = 3$, $n_{\min} = 25$). Bicubic and Lanczos interpolation methods yield nearly identical results. The dashed red line indicates the theoretical maximum efficiency, corresponding to perfect detection of all ground truth clusters with no false positives. Stricter clustering parameters (bottom rows) yield higher detection rates but identify fewer total clusters, reducing potential time savings.

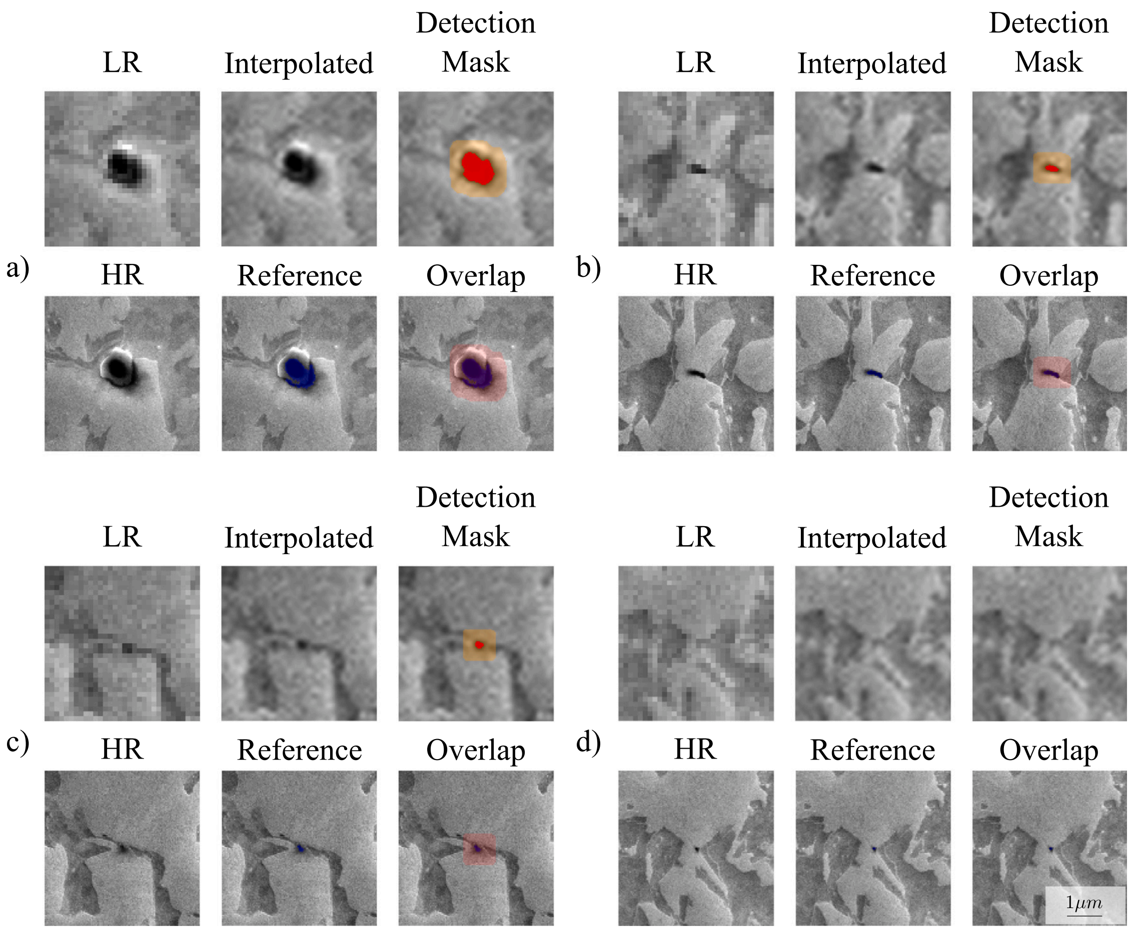

Fig. B.2. Representative damage sites at the 90 % operating point ($\tau = 60$, $n_{min} = 5$, $\epsilon = 1$), illustrating the four matching categories. Each panel shows, from left to right, the fast scan (θ_{fast}), the interpolated fast scan, the interpolated fast scan after thresholding, clustering, and dilation (detected clusters), the high-resolution analysis scan (θ_{analysis}), the analysis scan after thresholding and clustering (reference clusters), and the resulting overlap. Detected clusters in the interpolated fast scan are outlined in red, their dilated extent in orange, and the reference clusters from the analysis scan in blue. (a) Merge case: a single detected cluster in the fast scan overlaps two separate reference clusters, both of which are counted as detected. (b) A large damage site (reference cluster ≥30 pixel) correctly detected. (c) A small damage site (reference cluster <30 pixel) correctly detected. (d) A small damage site missed in the fast scan, with no overlapping detection. All crops use the $128\times \times 128$ pixel window used for evaluation.

Data availability

Data will be made available on request.

References

[1] C.C. Tasan, M. Diehl, D. Yan, M. Bechtold, F. Roters, L. Schemmann, C. Zheng, N. Peranio, D. Ponge, M. Koyama, K. Tsuzaki, D. Raabe, An overview of dual-phase steels: Advances in microstructure-oriented processing and micromechanically guided design, Annu. Rev. Mater. Res. 45 (Volume 45, 2015) (2015) 391–431, http://dx.doi.org/10.1146/annurev-matsci-070214-021103.

[2] C. Kusche, T. Reclik, M. Freund, T. Al-Samman, U. Kerzel, S. Korte-Kerzel, Large-area, high-resolution characterisation and classification of damage mechanisms in dual-phase steel using deep learning, PLoS One 14 (5) (2019) e0216493, http://dx.doi.org/10.1371/journal.pone.0216493.

[3] H.V. Atkinson, G. Shi, Characterization of inclusions in clean steels: a review including the statistics of extremes methods, Prog. Mater. Sci. 48 (5) (2003) 457–520, http://dx.doi.org/10.1016/S0079-6425(02)00014-2.

[4] M.A. Wollenweber, C.F. Kusche, T. Al-Samman, S. Korte-Kerzel, On the automated characterisation of inclusion-induced damage in 16MnCrS5 case-hardening steel, Adv. Ind. Manuf. Eng. 7 (2023) 100123, http://dx.doi.org/10.1016/j.aime.2023.100123.

[5] Y. Ju, S. Li, X. Yuan, L. Cui, A. Godfrey, Y. Yan, Z. Cheng, X. Zhong, J. Zhu, A macro-nano-atomic–scale high-throughput approach for material research, Sci. Adv. 7 (49) (2021) eabj8804, http://dx.doi.org/10.1126/sciadv.abj8804.

[6] S. Medghalchi, C.F. Kusche, E. Karimi, U. Kerzel, S. Korte-Kerzel, Damage analysis in dual-phase steel using deep learning: Transfer from uniaxial to biaxial straining conditions by image data augmentation, JOM 72 (12) (2020) 4420–4430, http://dx.doi.org/10.1007/s11837-020-04404-0.

[7] A. Eberle, S. Mikula, R. Schalek, J. Lichtman, M.K. Tate, D. Zeidler, High-resolution, high-throughput imaging with a multibeam scanning electron microscope, J. Microsc. 259 (2) (2015) 114–120, http://dx.doi.org/10.1111/jmi.12224.

[8] A.L. Eberle, D. Zeidler, Multi-beam scanning electron microscopy for high-throughput imaging in connectomics research, Front. Neuroanat. 12 (2018) http://dx.doi.org/10.3389/fnana.2018.00112.

[9] D. Yan, C.C. Tasan, D. Raabe, High resolution in situ mapping of microstrain and microstructure evolution reveals damage resistance criteria in dual phase steels, Acta Mater. 96 (2015) 399–409, http://dx.doi.org/10.1016/j.actamat.2015.05.038.

[10] C.C. Tasan, J.P.M. Hoefnagels, M. Diehl, D. Yan, F. Roters, D. Raabe, Strain localization and damage in dual phase steels investigated by coupled *in-situ* deformation experiments and crystal plasticity simulations, Int. J. Plast. 63 (2014) 198–210, http://dx.doi.org/10.1016/j.ijplas.2014.06.004.

[11] S.R. Babu, S.K. Michelic, Overview of application of automated SEM/EDS measurements for inclusion characterization in steelmaking, MetalMat 1 (2) (2024) e18, http://dx.doi.org/10.1002/metm.18.

[12] T. Merryman, J. Kovacevic, An adaptive multirate algorithm for acquisition of fluorescence microscopy data sets, IEEE Trans. Image Process. 14 (9) (2005) 1246–1253, http://dx.doi.org/10.1109/TIP.2005.855861.

[13] V.S. Tong, A.J. Knowles, D. Dye, T.B. Britton, Rapid electron backscatter diffraction mapping: Painting by numbers, Mater. Charact. 147 (2019) 271–279, http://dx.doi.org/10.1016/j.matchar.2018.11.014.

[14] G.M.D.P. Godaliyadda, D.H. Ye, M.D. Uchic, M.A. Groeber, G.T. Buzzard, C.A. Bouman, A framework for dynamic image sampling based on supervised learning, IEEE Trans. Comput. Imaging 4 (1) (2018) 1–16, http://dx.doi.org/10.1109/TCI.2017.2777482.

[15] Y. Zhang, G.M.D. Godaliyadda, N. Ferrier, E.B. Gulsoy, C.A. Bouman, C. Phatak, G.M.D. Godaliyadda, N. Ferrier, E.B. Gulsoy, C.A. Bouman, C. Phatak, SLADS-Net: Supervised learning approach for dynamic sampling using deep neural networks, Electron. Imaging 30 (2018) 1–1316, http://dx.doi.org/10.2352/ISSN.2470-1173.2018.15.COIMG-131.

[16] T. Dahmen, M. Engstler, C. Pauly, P. Trampert, N. de Jonge, F. Mücklich, P. Slusallek, Feature adaptive sampling for scanning electron microscopy, Sci. Rep. 6 (1) (2016) 25350, http://dx.doi.org/10.1038/srep25350.

[17] S. Clusiau, N. Piché, N. Brodusch, M. Strauss, R. Gauvin, Workflow automation of SEM acquisitions and feature tracking, Ultramicroscopy 269 (2025) 114093, http://dx.doi.org/10.1016/j.ultramic.2024.114093.

[18] A.E. Tekkaya, N. Ben Khalifa, O. Hering, R. Meya, S. Myslicki, F. Walther, Forming-induced damage and its effects on product properties, CIRP Ann 66 (1) (2017) 281–284, http://dx.doi.org/10.1016/j.cirp.2017.04.113.

[19] O. Hering, A.E. Tekkaya, Damage-induced performance variations of cold forged parts, J. Mater. Process. Technol. 279 (2020) 116556, http://dx.doi.org/10.1016/j.jmatprotec.2019.116556.

[20] F. Hannard, T. Pardoen, E. Maire, C. Le Bourlot, R. Mokso, A. Simar, Characterization and micromechanical modelling of microstructural heterogeneity effects on ductile fracture of 6xxx aluminium alloys, Acta Mater. 103 (2016) 558–572, http://dx.doi.org/10.1016/j.actamat.2015.10.008.

[21] G. Avramovic-Cingara, Y. Ososkov, M.K. Jain, D.S. Wilkinson, Effect of martensite distribution on damage behaviour in DP600 dual phase steels, Mater. Sci. Eng.: A 516 (1) (2009) 7–16, http://dx.doi.org/10.1016/j.msea.2009.03.055.

[22] E.E. Aşık, E.S. Perdahcıoğlu, A.H. van den Boogaard, Microscopic investigation of damage mechanisms and anisotropic evolution of damage in DP600, Mater. Sci. Eng.: A 739 (2019) 348–356, http://dx.doi.org/10.1016/j.msea.2018.10.018.

[23] L. Zhou, H. Wen, I.C. Kuschnerus, S.L.Y. Chang, Efficientand robust automated segmentation of nanoparticles and aggregates from transmission electron microscopy images with highly complex backgrounds, Nanomaterials 14 (14) (2024) 1169, http://dx.doi.org/10.3390/nano14141169.

[24] A. Shah, J.A. Schiller, I. Ramos, J. Serrano, D.K. Adams, S. Tawfick, E. Ertekin, Automated image segmentation of scanning electron microscopy images of graphene using U-net neural network, Mater. Today Commun. 35 (2023) 106127, http://dx.doi.org/10.1016/j.mtcomm.2023.106127.

[25] M. Rocklin, Dask: Parallel computation with blocked algorithms and task scheduling, in: SciPy 2015, 2015, http://dx.doi.org/10.25080/Majora-7b98e3ed-013.

[26] F. Kolpak, H. Traphöner, O. Hering, A.E. Tekkaya, Large strain flow curves of sheet metals by sheet extrusion, CIRP Ann 70 (1) (2021) 247–250, http://dx.doi.org/10.1016/j.cirp.2021.03.023.

[27] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, A density-based algorithm for discovering clusters in large spatial databases with noise, in: Proceedings of the Second International Conference on Knowledge Discovery and Data Mining, KDD '96, AAAI Press, Portland, Oregon, 1996, pp. 226–231.

[28] C.E. Duchon, Lanczos filtering in one and two dimensions, J. Appl. Meteorol. Clim. 18 (8) (1979) 1016–1022, http://dx.doi.org/10.1175/1520-0450(1979)018<1016:LFIOAT>2.0.CO;2.

[29] L.R. Dice, Measures of the amount of ecologic association between species, Ecology 26 (3) (1945) 297–302, http://dx.doi.org/10.2307/1932409.

[30] T. Reclik, S. Medghalchi, P. Schumacher, M.A. Wollenweber, T. Al-Samman, S. Korte-Kerzel, U. Kerzel, Resolution enhancement of scanning electron micrographs using artificial intelligence, Mater. Des. 253 (2025) 113955, http://dx.doi.org/10.1016/j.matdes.2025.113955.

[31] S. Medghalchi, E. Karimi, S.-H. Lee, B. Berkels, U. Kerzel, S. Korte-Kerzel, Three-dimensional characterisation of deformation-induced damage in dual phase steel using deep learning, Mater. Des. 232 (2023) 112108, http://dx.doi.org/10.1016/j.matdes.2023.112108.

[32] S. Medghalchi, J. Kortmann, S.-H. Lee, E. Karimi, U. Kerzel, S. Korte-Kerzel, Automated segmentation of large image datasets using artificial intelligence for microstructure characterisation and damage analysis, Mater. Des. 243 (2024) 113031, http://dx.doi.org/10.1016/j.matdes.2024.113031.