

Chapter 10

Robotic Inspection Systems

Christian Eitzinger, Sebastian Zambal and Petra Thanner

Abstract Industrial quality control often includes the inspection of parts of complex geometry. While such an inspection can be quite easily done by humans, it poses certain challenges if the task is to be automated. Quite often, robots are used for handling the part to acquire a large number of images, each showing a certain area of the surface. The process of acquiring sequences of multiple images also has implications for the machine vision and analysis methods used in such tasks. This chapter covers all topics that relate to the implementation of robotic inspection systems for industrial quality control. The focus is on machine vision, while aspects that deal with robotics will only be addressed at a conceptual level.

Contents

10.1	Introduction	322
10.2	General Concept and System Overview	324
10.2.1	Offline Components	325
10.2.2	Online Components	328
10.3	Carbon Fibre Part Inspection	330
10.3.1	Image Acquisition Set-up	332
10.3.2	Process Model	335
10.3.3	Image Analysis	337
10.3.4	Prototype Implementation	338
10.4	Thermographic Crack Detection	338
10.4.1	Image Acquisition Set-up	341
10.4.2	Process Model	343
10.4.3	Image Analysis	344
10.4.4	Prototype Implementation	347
10.5	Conclusions and Outlook	348
	References	348

C. Eitzinger (✉) · S. Zambal · P. Thanner
Profactor GmbH, Im Stadgut A2, 4407 Steyr-Gleink, Austria
e-mail: christian.eitzinger@profactor.at

10.1 Introduction

Industrial quality control is often realized as an automated end-of-line inspection system that checks certain features on the products. Inspection covers functional aspects of the product, such as whether all components are mounted on an assembly, as well as aesthetic properties, such as checking for scratches on a polished surface. We would like to make a clear distinction between *inspection* and *measurement*. Measurement systems provide quantitative information about a product in the form of certain physical units that can be traced back to internationally accepted standards. Measurement systems usually aim at achieving high precision. Inspection systems generate qualitative information about a product, usually by assigning the product to a certain class, e.g. “accept”, “reject” or “rework”. The main property of an inspection system is its classification accuracy, often represented in the form of a confusion matrix.

Robotic inspection systems use robots to perform the inspection task. The term “inspection robot” is often used in conjunction with mobile robots used for the inspection of tunnels, channels or pipes [1], sometimes with highly specific kinematics [2]. In [3], a standard industrial robot is used for quality control of printed circuit boards. However, in this application, the robot is only used to handle the part and does not directly support the inspection process. The kind of inspection robots that we will be investigating in this chapter is based on standard industrial robots that guide a vision-based sensor system over a part of complex geometry, such as those described in Woern et al. [4]. In [5], the idea is extended with a focus on selectively reworking of defective areas on the part. The common aspect of visual inspection with robots is that multiple images of a single part are required, where each image provides an optimal view of an area to be inspected on the part. The key elements of optimizing the viewpoints of a vision sensor are covered in Yang and Ciarallo [6] and in the references mentioned there. The reason for choosing a robot over using multiple cameras—which is often the cheaper solution—is the higher flexibility of the robot and the possibility of adapting the inspection process to different product variants. Depending on the application, the robot may handle either the image acquisition system, including the camera and the illumination [7], or the product itself. Within the class of robotic inspection systems, we also want to distinguish between two different types of inspection tasks: in the first, simpler case, the robot is positioning the camera at a sequence of pre-defined positions and an image is acquired from that position. Typically, the number of positions is low, e.g. up to 15 different locations. Each image is processed separately and the processing steps and parameters are set up manually. A typical application is completeness inspection of a complex assembly. The inspection system has to check whether the assembly contains all components, whether the correct type of component is mounted on the assembly and whether the components are in their correct positions. An example is shown in Fig. 10.1.

A typical scenario is the inspection of a fully assembled car engine, where the correct mounting of several plugs, pumps, hoses and cables needs to be verified.




Fig. 10.1 An inspection robot for completeness inspection in a laboratory setting. The insert shows a point cloud of the part with the false component *highlighted in red*

Obviously, there is no single viewpoint from which all these components can be seen, and therefore, the camera needs to be positioned to obtain an optimal view of each component. In this case, the inspection task is—at a conceptual level—identical to standard machine vision systems. Single images are acquired, they are analysed independently, and a decision is made based on the results of each single image.

In this chapter, however, we want to deal with the second, more complex case, where the robot needs to inspect a larger area by moving a camera over the whole surface. In this case, a very large number of overlapping images are acquired, typically several hundreds of images. These images cannot be processed independently, and manually setting up the inspection process is impossible, simply due to the large number of images. The typical application is surface inspection, where defects have to be found on the surface of a part. One such scenario is the inspection of chrome-plated fittings for small scratches and dents. The shiny surface allows an inspection only for a specific geometrical arrangement of illumination, camera and surface, so that a 3D free-form surface can only be fully covered by splitting it into a large number of regions and finding an optimal viewpoint for each of these regions. Due to the large number of regions, the inspection process cannot be done in a stop-and-go motion, but needs to be done by continuously moving the camera across the surface and adjusting its position in all six degrees of freedom.

This kind of inspection process poses several challenges that will be addressed in the remaining parts of this chapter:

- The large number of images does not allow a manual set-up of the whole inspection process. Many of the parameterization steps need to be done automatically, such as identifying the regions in the image that are of sufficient quality to be used for processing.
- The robot's path cannot be generated by a teaching process, because it is impossible for a human teacher to maintain an accurate distance and angle to the local surface normal. Instead, a process model of the image acquisition process is required that provides the basis for automatic planning of the robot's path.
- In many advanced inspection processes, several images of the same location are required, for example with different shutter times to extend the dynamic range of the images. To merge such images, accurate registration is needed, which in many cases requires fairly accurate synchronization between the robot's motion and the image acquisition process.

In the following sections, we will first describe how these challenges are addressed at a conceptual level and then proceed by presenting the details of two different, typical realizations of this concept.

10.2 General Concept and System Overview

A robotic inspection system is a complex technical system that requires a close interaction between its single components. The initial set-up of an inspection task for a 3D free-form surface cannot be done manually, because a human operator cannot define with sufficient accuracy the path that the robot has to follow. In many cases, the set-up of the image acquisition system relies on a specific geometrical configuration between camera, illumination and surface. This configuration only allows minimal tolerances, and consequently, manual positioning will not be an option. Therefore, robotic inspection systems need a (semi-)automatic offline configuration process that is used to define the whole inspection task so that it can afterwards be executed fully automatically. This includes elements such as 3D CAD models of the part, a process model of the image acquisition process, algorithms for path planning and calibration methods. All of these components will be further described in Sect. 10.2.1.

Once the offline process is finished and the inspection task has been defined, the robot should be able to automatically execute the inspection process, which may also include compensation for small deviations that may occur, e.g. in the positioning of the product to be inspected. This online inspection process includes the standard machine vision elements, such as low-level image processing, image segmentation, feature extraction, classification and decision-making. For robotic inspection systems, additional elements are needed, such as synchronization

between image acquisition and the robot’s position, projection of defects from the 2D images to the 3D surface and identifying areas in the image that are suitable for inspection. These topics will be addressed in Sect. 10.2.2.

10.2.1 Offline Components

Before the automatic inspection can be started, the inspection process has to be set up for the particular product variant. While in many machine vision tasks this set-up can be done manually, the robotic inspection tasks that we will discuss in the following sections require a semi-automatic set-up process because of the complexity of the task. This process is shown in Fig. 10.2.

The usual starting point is a 3D CAD model of the object to be inspected. This model needs to be stripped of all elements that do not include shape information such as measurements and other metadata. In many cases, the CAD model is converted into a simple standard format to remove the complexity of 3D free-form shapes that may be present. This can be achieved either by point clouds or preferably by a triangular mesh. However, shape information alone is not sufficient, and it needs to be augmented with additional information that is not typically included in CAD formats [8].

Such information may include, e.g., the inspection parameters, such as allowable defect sizes per defect type or the acceptable maximum number of defects. These parameters may be different for different areas on the product. Therefore, the CAD model sometimes needs to be split into areas to which different quality criteria apply. Occasionally, additional information will also include surface properties, such as whether the surface is milled, sand-blasted or polished. This can be used to switch image acquisition parameters during the acquisition process.

Once these augmented CAD data are available, we can use them to generate a path for the robot. This path has to fulfil several criteria:

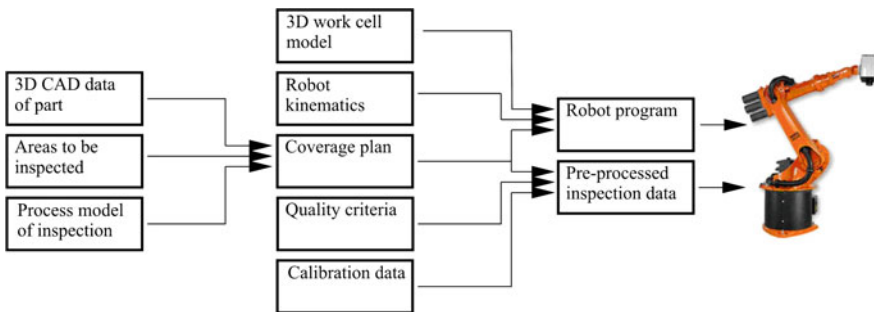


Fig. 10.2 Offline processing: the components and their interaction

- The path has to cover all surface elements that need to be inspected. For many inspection tasks, several images of a single surface element need to be acquired, e.g. under different illumination or with different shutter times. Therefore, each surface element has to be acquired multiple times, while the robot is moving.
- The path has to comply with certain time constraints, e.g. the speed may not be so fast that the maximum frame rate of the camera is exceeded, or if the frame rate cannot be adapted, the robot's speed has to be kept within quite tight tolerances.
- All the positions along the path that the robot is expected to follow must be reachable for the robot, i.e. they must be within the workspace of the robot.
- The whole path of the robot must be free of collisions, i.e. the robot shall not hit the product, itself or any other component that is inside the workcell.

These criteria have to be jointly fulfilled and constitute the boundary conditions of the path planning problem. In order to identify a single solution for the path, an optimization process is implemented to minimize the time needed for completing the whole path. In order to make sure that the robots smoothly follow the path and do not permanently change its configuration, it usually pays off to split the whole part into different zones that are treated sequentially. Quite often, such zones can be geometrical primitives such as (parts of) cylinders, which enable the use of pre-defined strategies for path planning.

Path planning has been investigated for many camera-based inspection problems, such as completeness inspection [9], for automated 3D object reconstruction [10] or dimensional measurement using laser scanners [11]. For such camera-based processes, it is also known as “view planning problem”. Usually, the length of the path or the number of views is chosen as an optimization criterion. In order to solve the associated optimization problem, algorithms that can deal with discrete search spaces, such as genetic algorithms [12] or particle swarm optimization, are regularly used. The basic problem structure is also known as the “art gallery problem” or the “watchman route problem”.

Path planning is a computationally hard problem [13], so that even with a low number of waypoints finding a globally optimal path often is computationally infeasible. Moreover, a globally optimal path with focus on path length is not always desirable because other attributes such as the ability for dynamic replanning can also be important. Recently, large neighbourhood search [14] has shown remarkable results for path planning in constrained environments. Also, ant colony optimization [15] has been successfully used to solve large-scale route planning problems with hard constraints [16]. We will not go into further detail on how the path is planned and converted into a robot program. We refer the reader to the literature above and focus our discussion only on the machine vision aspects of the path planning.

At the heart of the path planning, algorithm is a process model that represents the image acquisition process. This process model answers the question, which area on the part can be inspected from a given viewpoint. It includes not only basic elements such as the optical properties (focal plane, field of view, etc.), but also more complex ones such as the range of object distances that are still in focus and that

deliver images of sufficient quality so that they are suitable for inspection. Some of these parameters can only be determined by experiments. The process model is highly specific for the particular image acquisition set-up that is used. In the following sections, we will provide two concrete examples of such process models.

The second major offline step is the calibration of the whole inspection system. Calibration provides information about the relative position and orientation of all the coordinate systems that are involved in the inspection task. There are multiple coordinate systems that we will discuss in the following paragraphs. The first one is the world coordinate system. This system is usually linked to some significant physical object in the workcell, whose position can be determined with sufficient accuracy. The world coordinate system is often used for simulations, where all objects in the robotic workcell are represented and need to be placed in 3D space relative to that coordinate system. From the world coordinate system, we proceed to the tool centre point of the robot. The tool centre point is linked to the world coordinate system via a kinematic chain across all the joints of the robot. This kinematic chain and the coordinate transformations associated with it are assumed to be known with sufficient accuracy. They are defined by the type of robot and are part of the specification of the robot. The position of the tool centre point relative to the world coordinate system is calibrated by making the robot point to a few particular, well-defined positions in the world coordinate system. Depending on the accuracy that is required, a pointing tool, such as a needle or stylus, may be used. The calibration procedure is often semi-automatic, where the needle is manually positioned and the transformations are calculated automatically. This functionality is typically a part of the robot's control software and does not need any further consideration when setting up a robotic inspection system. Assuming that the camera is mounted on the robot's gripper, a calibration is needed between the tool centre point and the camera coordinate system. This is called hand-eye calibration, and many different methods have so far been developed for this purpose [17–19]. Camera calibration involves two steps that may be executed separately or jointly: the intrinsic calibration and the extrinsic calibration. The intrinsic calibration is used to determine all parameters that belong to the optical system of the camera, such as the focal length, the centre of the optical axis, scaling and various parameters related to optical distortions. The extrinsic calibration determines the position of the camera, i.e. the centre of the optical axis, relative to other coordinate systems, in our case relative to the tool centre point. Intrinsic calibration and extrinsic calibration are done by placing a calibration object in front of the camera in different positions (or by moving the robot) and acquiring a series of images. Quite often, the calibration object includes a checkerboard pattern [20, 21], whose corners can be accurately localized in the images. The coordinates of these corners are assumed to be known and are then set in relation to the pixel coordinates of the camera, so that the intrinsic and extrinsic camera parameters can be determined.

The last coordinate system that we need to consider is the workpiece coordinate system. Usually, the 3D shape of the product to be inspected is defined in a 3D CAD model that uses a particular coordinate system that is linked with certain features on the product. In order to inspect the product, the inspection system needs to know

where the product is placed in space. Quite often, it is not needed to determine the position of each single product to be inspected, because manufacturing tolerances are sufficiently low so that the products can be considered as being identical for the purpose of inspection. Therefore, the standard solution is to have a mechanical fixture in which the product is placed and accurately aligned relative to the fixture. By having a set of fiducial marks on the fixture, it can be accurately positioned relative to the robot and the world coordinate system. In many cases, the robot's control unit also supports the definition of a product coordinate system, e.g. by again using a needle or stylus to identify the position of the marks relative to the robot coordinate system. Alternatively, the camera may be used to localize the marks.

Once we performed all these calibration steps and all the coordinate systems are determined to the required level of accuracy, we are able to map defects that were found in the image onto the product in 3D space. Starting from the pixel coordinates of the defect in the camera image, we use the intrinsic calibration and extrinsic calibration to transform the position into the tool coordinate system. By recording all the joint angles at the time the image was taken, we can use the kinematic chain of the robot to transform the position into the world coordinate system and with a final transformation into the product coordinate system. Actually, it is not as simple as that, because the camera only provides 2D information that represents a straight line in 3D space on which the defect is localized. This line then needs to be intersected with the 3D model of the product to determine the location of the defect on the product, but in any case, the sequence of coordinate transformations is the same as described before.

The final element that has to be preset before the automatic inspection can commence is the quality criterion that has to be applied. These criteria convert the list of defects into a single “accept”/“reject” classification. The criteria will be different for different areas on the part. For regions, such as sealing areas, tighter criteria will apply than for regions that do not fulfil a particular function. Initially, these criteria are specific in 3D, i.e. based on the CAD drawing of the part. In order to save computing time during online inspection, it is recommendable to pre-process the regions and compute back-projections of these regions into the single images taken from the different viewpoints. It is then much simpler and less time-consuming to identify those areas in the image that have to be inspected.

10.2.2 Online Components

After setting up the inspection parameters, the robot's path and all the coordinate systems, the robotic inspection system is ready for operation. During the execution of the inspection process, a set of distinct modules are operating in parallel that are described in the paragraphs below. The description will follow the logical processing steps from the image to the high-level decision-making as shown in Fig. 10.3 and will highlight the specific properties of these modules.




Fig. 10.3 Data flow during online processing

Processing starts with the images. Depending on the application, this may be just a single image or a set of images taken under different illumination. For an efficient implementation of the inspection process, we will avoid stop-and-go motion for each image and instead assume a continuous scanning process. This implies that all the images are taken at slightly different positions, which needs to be taken into account when the joint analysis of several images, e.g., with different illumination is needed for the application. For that kind of analysis, it is necessary to determine the exact location of a single, particular point on the part’s surface in the set of images. Due to the continuous motion, this point will appear in different locations in the images. To facilitate image registration, accurate position information is required that usually comes from the robot’s control unit. A common approach is to synchronize image acquisition and the robot’s motion either electrically, where the camera provides an electronic trigger signal whenever an image is acquired, or by software, where time stamps are used to perform the synchronization. Recently, communication protocols such as the precision time protocol (PTP) have been developed to solve this problem. The process of merging image data and robot position information is done by a module that we would like to call data assembler. The task of this module is to create a sequence of position-referenced images, so that each image also contains information about the joint angles of the robot at the time of acquisition. The cycle times of the camera (e.g. 200 frames per second) may be different from the cycle times of the position read-out module (e.g. 10 ms), which may require interpolation between time instances.

The sequence of position-referenced images is then fed into an image processing module. The exact processing that takes place totally depends on the application. Examples of such processing will be described in the following sections. Common to most of the applications is the need to process comparably large amounts of data. This is best done by a pipeline-based architecture, where the single processing steps are done in parallel for multiple images taken at different times. This processing pipeline will ensure that the amount of data can be processed in time; however, it has to be noted that there will always be a delay between the time the image is taken and the time by which the processing result is available. However, in robotic inspection systems, where we need to wait until the inspection of the whole part is finished, this will not be a problem. Typically, we expect delays of about 1 s, while inspection of the whole part takes 30–60 s.

The sequence of processing steps follows a typical scheme that is widely used in surface inspection tasks. The first step is the image segmentation. This step is highly

application specific and will look totally different, depending on the surface texture, the reflection properties and on the physical processes (e.g. in thermography) that are guiding the image acquisition. Goal of the segmentation is to extract a set of image patches that show “interesting” areas. We call these “areas of interest” or “events”, because the fact that it is detected by image segmentation does not necessarily imply that it is a real defect. To determine whether there is a real defect, the image patches are converted into feature vectors. Again, the features may be application dependent; however, at this processing level, quite generic features exist that are useful for a very wide range of inspection tasks. The feature vectors are then fed into a classification system that determines whether there is a defect, and if so, the particular type of defect is shown. Features used in this classification may include the grey-level distribution inside of the defect (texture) or depth information coming from photometric stereo calculations. Sometimes, calculation of these features can be optimized to improve classification accuracy [22]. The main challenge is to distinguish the real defects from the many other things that may appear on the surface.

The final step after classifying the different types of defects is an accept/reject decision for the whole part. There are two substantially different cases. In the first—more seldom—case, the decision is based on the subjective judgment of an expert doing quality control. This happens, when aesthetics have to be judged or other properties that are difficult to assess in an objective manner. In order to convert this expert’s knowledge into an automatic decision, machine learning methods are used that extract decision rules from a set of samples [23]. In the second case, the decision is based on quality control criteria that are documented in a quality management system. This criteria are written down in the form of rules including specific thresholds that determine limits for properties such as the maximum size of a defect, the minimal allowed distance between defects, the maximum number of defects and other rules that deal, e.g., with clusters of defects. Quite often, however, we observe differences between the actual manual inspection process and the quality criteria. Human experts who performed the inspection often gather specific knowledge about the production processes and adjust their decision-making based on their observations [24]. These adjustments are typically not documented in the quality management system. In order to make sure that the inspection robot best reproduces the actual (human) inspection process, statistical analysis of samples and machine learning will be applied here as well.

10.3 Carbon Fibre Part Inspection

Composite materials are nowadays gaining more and more importance in many branches of industry. Classically, composite materials were limited to very special applications such as aerospace or high-end sports equipment. However, the use of these materials is getting more and more popular for many different applications. The widespread use in a large spectrum of products increases the need for

production processes with a high degree of automation and reduced costs. In this context, automated inspection and process control based on machine vision are very fruitful technologies.

In the following, we focus on surface inspection of carbon fibre and glass fibre-reinforced plastics. These materials exhibit rather complex optical properties as they are shiny and provide either too little or too high contrast. This results in too dark or too bright patches in the images and prevents further processing. With respect to the design of a machine vision system, the following properties of fibre-reinforced plastics are thus relevant:

- The mirror-like specular reflection (carbon fibre and glass fibre behave like tiny cylindrical mirrors) puts several challenges on a reliable machine vision system.
- Raw textile fibre materials typically exhibit a repeating pattern (e.g. woven material, sewing yarns). Such repeating patterns make it difficult to stitch images by means of computer vision methods [25, 26].

Regarding surface errors that need to be detected by a surface inspection system, two relevant scenarios are identified:

- Inline inspection for automated post-treatment.
- Inspection for final quality control and good/bad decision.

A crucial property of fibre-reinforced plastics is the orientation of fibres within the material. Fibre orientation matters because of two different reasons. First, mechanical properties depend strongly on the orientation of the fibres. A structural carbon fibre part can only deal with strong forces that act upon it, if forces and carbon fibres are well aligned. This is often designed using finite element calculations, and much effort is spent on minimizing the weight for given mechanical properties and loads. Second, for those parts that are visible in the final product, fibre orientation matters because of aesthetic reasons. Typically, fibres should be aligned with the edges of the part, and distortions of the woven material are not allowed.

In the remainder of this section, we describe a robotic surface inspection system that inspects fibre angles on the surface of 3D parts with moderate geometric complexity. The system is able to determine the fibre angles on the surface of carbon fibre or glass fibre-reinforced plastics, fully covering the part. Of course, there are some physical limitations, e.g., in tight corners of the parts, in areas that cannot be reached by the robot and in places where collisions would occur.

We focus on the surfaces of so-called pre-forms. In the production process of fibre-reinforced plastics (FRP), pre-forms represent an intermediate step. Basically, a pre-form consists of fibre textiles which are already put roughly into the shape of the final part. In subsequent production steps, resin is injected into the part or the pre-impregnated pre-form is heated. In both cases, the liquid resin tightly fills all the space between the fibres. After curing, the part keeps its shape on its own. A critical point is that pre-forms are easily deformable. In this state of the production process, the fibres are not kept in place by the surrounding plastic matrix. As a result, the actual shape of the part does not perfectly match the CAD data, which requires us to first obtain sufficiently accurate 3D data. Hence, we propose a separate geometry capturing

step using a laser range scanner or low-cost alternatives such as RGB-D sensors which are gaining more and more popularity for 3D object reconstruction [27].

For the inspection of pre-forms, it is in general more suitable to mount the camera on the robot and move it over the part. Keeping the camera fixed and moving the pre-form part is not an option as the pre-form would deform and could even be damaged. The presented system supports two modes: stop-and-go inspection, where the robot is stopped for capturing a set of images and continuous capturing of images. While continuous capturing of images is the far more interesting approach, we will also discuss aspects of stop-and-go inspection.

10.3.1 Image Acquisition Set-up

There exist different approaches to perform automated inspection of carbon fibre-reinforced plastics. We shall ignore general non-destructive testing methods such as computed tomography, ultrasound and the like and focus on machine vision. One possible approach to inspection using machine vision is to suppress specular reflections of carbon fibre and glass fibre. This may be tackled with diffuse illumination of the inspected surfaces [28]. The resulting image ideally does not show any specular reflections and is homogeneous, occasionally with low contrast. Texture analysis methods can then be used to calculate fibre angles or perform defect detection. However, this method proved to be somewhat sensitive to changes in the illumination and to the properties of the material that is inspected. Also there are difficulties when inspecting clear-coated parts. We are thus following a different strategy for the described fibre angle measurement system. Instead of suppressing the specular nature of fibres, we use a reflection model that exploits the specular reflectance of the fibres [29, 30]. Based on a set of images that are taken under different illumination conditions, the reflection model makes it possible to calculate fibre orientations. Basically, a special form of photometric stereo [31, 32] is applied.

Carbon fibre and glass fibre are modelled as very thin cylindrical mirrors. A light ray that hits a small cylindrical mirror is reflected in the shape of a cone. Typically, carbon fibres have a diameter of approximately $5\text{--}9\ \mu$, while camera-based surface inspection systems cover a surface region of at least $20 \times 20\ \mu\text{m}^2$ per pixel. Hence, the corresponding light rays have a sufficiently large diameter, so that the cylinder may be considered as being infinitely small and the cone-shaped reflection model remains valid.

Given a fixed light position, the exact shape and orientation of the reflection cone depend on the orientation of the reflecting fibre. Inversely, given the parameters of the reflection cone, it is possible to determine the orientation of the fibre in 3D space. By capturing images illuminated from different directions, it is possible to determine the reflection cone parameters and, hence, also the fibre orientation on the surface.

For the following considerations, we concentrate on the inverse direction of light. We follow the line of sight from the camera's optical centre (denoted c in Fig. 10.4) to the intersection with a fibre (denoted O). According to the




Fig. 10.4 Fibre reflection analysis (from [29])

considerations above, this ray is also “reflected” in the shape of a cone. Light sources are located on a circle with centre l_0 . The intersections of the light circle and the reflection cone are marked as positions l_1 and l_2 . In an ideal model, only light from these two positions on the circle is reflected into the camera.

By capturing many images with different light sources switched on, the intersection of the reflection cone of the line of sight with the circle of illuminations is calculated. This is done by investigating the variation of brightness of pixels corresponding to a single point on the surface. In case of point light sources, theoretically, an infinite number of light sources on the circle are required in order to exactly hit the intersection points. The diagram on the right top of Fig. 10.4 shows the variation of grey values for corresponding pixels with point light sources. The dotted line shows the intensity of reflections for an infinite dense number of light sources. The continuous line shows the intensities for a limited number of 8 light sources. The peaks of the dotted line indicate the intersection points of light circle and reflection cone. Knowing these intersection points, it is easy to calculate two normal vectors s_1 and s_2 of the fibre that points into direction f . Finally, the cross-product of s_1 and s_2 is equal to f .

A dense distribution of light sources over the light ring is theoretically necessary. In order to solve this problem, broader segments of light sources are used. A set of 8 light segments are sufficient to reconstruct the distribution of grey values (Fig. 10.4, right bottom).

The approach that has been described above works well for static inspection, where the camera remains in the same place until all the images under different illumination have been acquired. If the camera is moving during the acquisition




Fig. 10.5 Image format for continuous capturing during motion—large format (*Left*) versus slim format (*Right*)

process, accurate registration between the images is needed. This raises a few open questions that will be addressed in the next paragraphs. First, we need to investigate the optimal size of images that should be captured. Many industrial cameras offer the capability of acquiring only a subset of all image pixels at the benefit of higher frame rates. Here, we consider two possible strategies: large (slow) versus slim (fast) image format. When capturing large images during motion, the individual raw images are captured at quite large distances (Fig. 10.5, left). Using a slim image format, images are captured at relatively short distances (Fig. 10.5, right). Note that the overlap between consecutive images of about 90 % is identical in both cases. Three issues need to be considered regarding the image format:

1. Inaccuracy in registration: in order to perform the previously described fibre reflection analysis, the very same surface point needs to be identified in a set of sequentially captured images. For a typical set-up with 8 raw images for fibre reflection analysis, a point needs to be visible in all 8 raw images. If the captured images are rather large, the distance between two successive images has to be large as well in order to minimize the time needed for inspection and avoid unnecessary overlaps. If the images are captured at large intervals, the large distance of two successive capture positions makes it difficult to establish accurate pixel-to-pixel correspondences.
2. Cycloid movement of light positions: in a static set-up, the surface is illuminated with light coming from different positions along a circle. If the camera is moving together with the light sources, the circle becomes a cycloid. If the temporal and spatial distance between two consecutive images is large, then this cycloid shape has to be considered in the analysis.
3. Curvature of surface transversal to motion: unlike the two previous points, this one concerns the size of the image in the direction that is orthogonal to the scanning direction. If the surface is strongly curved in this direction, the usable part of the captured image will be small.

Given these considerations, the general preference is to have a small image height in direction of the scanning motion and reasonably wide images orthogonal to the scanning direction. This simplifies image registration and optimizes the speed. Typical frame rates go up to 1000 images per second.

In order to assign the correct sensor position to the captured images, it is necessary to synchronize image acquisition with robot positions. We propose the use of time stamps to achieve this synchronization. An integrated clock is a standard feature in many modern GigE industrial cameras. This integrated clock is used to assign a time stamp to each image. Also, robot controllers usually contain a clock which we use to assign time stamps to continuously record robot positions. In fixed time intervals, the clocks of camera and robot controller are synchronized by calculation of the time offset. Clock synchronization makes it possible to correctly assign robot positions to the captured images. We are using a software implementation of this synchronization although hardware solutions (e.g. PTP hardware implementation) are also available and provide even higher accuracy.

10.3.2 Process Model

The basis for the process model of image formation is the fibre reflection analysis described in the previous section. The critical question which the process model has to answer is as follows: for which regions does the fibre reflection analysis work, given the position of the sensor relative to the surface. Of course, this includes the question which regions are basically visible by the camera system. This covers aspects such as the camera's field of view and in-focus depth. However, additional aspects need to be considered as well: in order to make fibre reflection analysis possible, it is required that the investigated fibres' orientations are within certain limits. Figure 10.6 shows typical results of fibre reflection analysis. In the leftmost image, the fibre angles are colour-coded, with red being horizontal (left-to-right) fibres and blue being vertical (top-to-bottom) fibres. The third image from left shows the same material wrapped around a cylinder. It is clearly visible that the analysis fails in regions where the surface is tilted with respect to the imaging plane. In case the reflection analysis fails, more or less random fibre angles are calculated, which corresponds to the noisy image regions.




Fig. 10.6 Fibre reflection analysis—flat surface (*Left*) versus cylindrical surface (*Right*)




Fig. 10.7 Illustration of fibre reflection for different combinations of in-plane and out-of-plane orientation. Fibre reflection analysis fails if no light from the LED ring reaches the camera (right most illustration)

For the cylindrical shape we see that not only the overall surface geometry influences the success of the analysis, but also the fibre orientation is relevant. Whereas for the horizontal (red) fibres the analysis fails already at low deflections ($\sim 12^\circ$ out of the imaging plane), correct calculations can still be done for (blue) fibres that run from top to bottom at angles of up to 40° . An explanation for this is illustrated in Fig. 10.7.

Based on this process model, we can describe at which fibre angles the analysis fails. However, at the point in time when offline path planning is made, the actual fibre orientations are not yet known. Even though (in-plane) fibre orientations are documented together with the CAD data (“ply book”), the fibre angles may deviate from these ideal orientations. The surface geometry that we know a priori somehow restricts fibre orientations, as it is reasonable to assume that fibres are almost parallel to the surface and out-of-plane angles are close to zero. However, for some materials (e.g. woven textiles) or defects that occur, the critical out-of-plane angle may be higher than expected and subsequently may lead to failure of fibre reflection analysis.

For some applications, it may be possible to safely estimate the range of expected out-of-plane angles. In this case, this range of angles is used in the offline path planning and failure of fibre reflection analysis is excluded. The drawback is that the estimated range of out-of-plane angles may be larger than actual out-of-plane angles. If this is the case, path planning may lead to a path that is much more complex than necessary. In order to avoid unnecessarily complex paths, dynamic replanning of inspection positions may be implemented. As soon as a failure of fibre reflection analysis is detected, the relevant surface region is captured from additional varying positions until a valid fibre angle can be calculated. The critical point is that dynamic replanning should not occur too often as it may introduce a significant increase of scanning time. Furthermore, it is not straightforward to implement automatic detection of failure of fibre angle calculation. One possible approach is to assess the noise level of calculated fibre orientations. If the noise level is too high for some image region, a failure is assumed. In general, the strategy for dynamic replanning and a reasonable range for expected out-of-plane fibre angles need to be adapted to the respective application. Relevant parameters are mainly the material used and the types of possible defects.




Fig. 10.8 The process model is used to calculate valid regions in the field of view. The *black* boundary marks the visible region of the camera. Valid regions (with respect to the process model) are *coloured green* (colour figure online)

Figure 10.8 shows different examples for results provided by the process model. Visible regions as seen from the inspection system are marked with black lines. The actually valid regions for which fibre reflection analysis is expected to work are restricted to much smaller regions coloured green.

10.3.3 Image Analysis

Once a full scan of a part is finished, fibre reflection analysis is performed using the captured image data from the 3D surface. Two approaches are possible:

1. Starting from the 3D CAD model, we identify the point for which we want to calculate the fibre orientation. For this surface point, the set of images needs to be determined that shows the point under different illuminations. By projecting the point into the images, we extract the single grey values and calculate the fibre orientation. This requires that the projection of each surface point in several images has to be calculated as it is not self-evident in which images the point will be visible. Nevertheless, this approach makes it easy to calculate fibre angles on a regular grid mapped to the 3D surface.
2. Fibre orientation can also be calculated in the 2D domain. Because the images stem from 3D surfaces, a mapping of surface points to 2D points is introduced and a transformation between consecutive images in 2D is calculated. For almost flat surfaces, this may be approximated with an affine transformation. For more complex surface geometries, also more general distortions need to be considered. The fact that multiple pixels (whole images) are transformed and analysed at once makes this approach very efficient. Once the fibre angles are calculated in 2D, they are projected back onto the 3D surface in order to obtain the fibre orientations on the inspected surface. With this approach, a dense mapping can be more easily achieved; however, the density on the 3D surface is not regular and will vary depending on the geometry.

Independent of the exact strategy for fibre angle calculation for individual surface points, the calibration accuracy is very critical. This is the case because an exact pixel-to-pixel mapping between individual raw images is desirable. For the described inspection system, a single pixel maps to a surface region with a diameter of approximately 20–80 μm . To establish an accuracy of positioning for pixel-to-pixel mapping at this scale is clearly a challenge. Of course, calibration of the whole inspection system including an accurate localization of the scanned part has to be done. In order to further increase the accuracy of image registration, image stitching methods that work on 3D surfaces [33] may be considered. The idea is to define an objective function that describes how well individual texture patches overlap. The objective function is optimized subject to parameters that describe the location of images projected onto the surface.

Most of the existing stitching methods aim at texturing of objects under diffuse illumination. In this context, it is comparatively easy to set up a reasonable objective function, e.g. based on a pixel-wise comparison of grey values. For fibre reflection analysis, the images are acquired under different illumination and images taken from the same spot may look very different. A pixel-based comparison of grey values is thus not likely to succeed. Instead, the noise level in the calculated fibre orientations may be used as an objective function for image alignment. General frequency measures of combined images such as those used in autofocus for cameras are also applicable.

10.3.4 Prototype Implementation

A prototype of a CFRP inspection system was implemented using a TX90L Stäubli industrial robot. The sensor system is based on a Genie TS 2500M industrial Gigabit Ethernet camera and a high-power LED light ring. A dedicated micro-controller triggers LED light sources and image acquisition of the camera. Different image formats, LED lighting sequences and other settings of the image acquisition system are supported and can be adapted to the respective application.

Figure 10.9 shows a picture of the sensor system together with a 3D visualization of the robot, the sensor viewing pyramid and the CAD model of the part. In the visualization, fibre orientations for a part of the CAD model are colour-coded and back-projected onto the CAD surface. Figure 10.10 shows the resulting fibre angles (colour-coded) for a region on a curved part.

10.4 Thermographic Crack Detection

Zero failure production is gaining more and more importance in industrial production processes for two main reasons: first in respect to reduce production cost due to minimizing waste material and second to ensure highest product quality for




Fig. 10.9 Prototype implementation: fibre orientation sensor scanning the surface of a part—picture (*Left*) and visualization (*Right*)




Fig. 10.10 Fibre angles mapped onto the surface of 3D-shaped part

the whole product life cycle time in order to avoid expensive callback actions. In particular, components working under strong mechanical and thermal stress need to be checked very carefully, since even small defects can affect performance and reliability in a negative way. Cracks are a common source of failure in metallic parts and thus require much effort for reliable crack checking.

The most common current procedure for crack detection for metallic parts is a process that dates back to the 1920s and is called “magnetic particle inspection” (MPI). This method is infamous in industry, because it is a cumbersome, dirty

process that is often done manually even in otherwise fully automatic production lines. The component to be tested is temporarily magnetized before applying a suspension of finely divided coloured or fluorescent magnetic particles. Cracks or inclusions disrupt the magnetic field and cause the magnetic particles to show up the crack. Using UV light under low-light conditions, the fluorescent particles in the suspension are activated to increase the visibility of the cracks. Afterwards, the component needs to be demagnetized. Magnetic particle inspection for complex parts is a manual process with well-known uncertainties in inspection performance. In combination with the necessity for 100 % inspection, an automated solution for crack detection is inevitable.

In the following, we focus on a non-destructive method based on infrared thermography which can be used for fully automated crack detection. Infrared thermography for non-destructive testing aims at the detection of (sub-) surface features, i.e. cracks or material anomalies, owing to temperature differences observed on the investigated surface during monitoring by an infrared or thermal camera [34]. An infrared camera detects and records the radiation emitted (electromagnetic radiation) by a material under investigation and converts this into a thermal image [35].

In general, infrared thermography is specified in two categories: passive and active. Passive thermography is used to investigate objects which are at a different temperature than ambient temperature. Most common applications for passive thermography are for thermal efficiency survey for buildings and infrastructure, predictive maintenance, medicine, fire detection and in non-destructive testing [36]. Active thermography uses an external heat source (i.e. flash lamps or laser) in order to induce a spatially and timely well-defined heat pulse to the component to be inspected causing a transient phase as long as heat distributes into the sound areas. This heat flux will be disturbed by voids such as surface or subsurface cracks. Fig. 10.11 illustrates the typical heat dissipation caused by a spatially limited heating. For homogenous materials, the heat dissipates uniform in all directions. Cracks disrupt the heat dissipation and cause non-uniform heat dissipation which can be detected by a thermal camera. Figure 10.11 illustrates the expectable heat dissipation caused by a point heat source and disruption by a surface near crack (right) and a crack at deeper location.




Fig. 10.11 Detection of subsurface cracks. The crack depth must be much smaller than its distance to the heat source

In the following, we focus on crack detection for metallic forged parts, more precisely for crankshafts. Crankshafts are available at a high number of different shapes and materials. They are parts of a high geometric complexity and possess different surface qualities. Typically, surface qualities of forged parts vary in the rang from rough unprocessed surfaces with typical emission coefficients of about 0,8 up to shiny processed surfaces with typical emission coefficient of 0,2. Basically, crack characteristics for the considered parts vary in a wide range in respect of position, orientation, size and depth. Cracks may also be closed at the product surface.

In the remainder of this section, we describe a robotic crack detection system which is able to inspect metallic forged parts of complex geometries using heat flux thermography. This system is able to perform a 100 % crack inspection of the sample parts. Of course, there are some physical limitations, e.g., in small gaps or areas that cannot be reached by camera and laser or regions where collisions may occur.

10.4.1 Image Acquisition Set-up

There exist several non-destructive testing methods to perform automated crack detection, for example eddy current, ultra sonic or X-ray. We focus on thermographic crack detection. One possible approach is to use passive thermography. In case of crack detection, passive thermography works only for surface-breaking cracks and materials with low-emission coefficients and requires the component to be tested to be at a significant higher temperature than the ambient. Caused by the different emission coefficients between material (low-emission coefficient) and crack (high emission coefficient), the crack occurs brighter than the surrounding material and can therefore easily be detected within the thermal image. Figure 10.12 shows two sample images acquired by using passive thermography methods. The left image shows a key with fatigue crack. The right image shows the front side of a gearwheel with a crack on one of its tooth.




Fig. 10.12 Sample images fro crack detection using passive thermography. The images show a key with fatigue crack (*Left*) and a gear-wheel with a crack on the front side of one tooth (*Right*)

However, this method proved to be pretty sensitive to surface qualities and does not perform for cracks which are closed or covered at the surface. We are thus focusing on active or transient thermography. In the case of active thermography, an external stimulus is necessary to induce relevant thermal contrast.

Inductive heating is a widely used and effective excitation method for crack detection of ferritic metal parts, e.g. for inline inspection of steel billets [37]. The technique uses induced eddy current to heat the material. Hence, inductive excitation is limited to be used for sample parts made from electric conductive materials and to detect cracks with known predominant crack directions. Since crack characteristics for crankshafts do not follow any predominant direction, we are thus following a more flexible strategy for heat induction by using laser spot heating.

Starting with a test part at thermal equilibrium, laser excitation induces a spatially and temporally well-defined heat impulse to the specimen. A spatially limited area (affected by the laser spot) will be heated compared to the sound material. This local temperature difference between spot position and sound area acts a motor for the heat flux. Heat starts to dissipate into the sound material. Heat flows until temporal equilibrium is reached. For homogenous materials, the heat flux is uniform and will be disrupted by any crack or void. The time period of the transient phase depends on the amount of induced energy (local temperature difference) and the heat conductivity of the material. Basically, metals are good heat conductors so the transient phase is temporal and spatially limited to a very short time period after excitation and area near the excitation position. By moving the laser spot, larger areas and hence whole test sample can be inspected. Figure 10.13 shows thermal images of a ring-shaped object with crack during laser excitation at different points of time (t_1 until t_6). The laser spot moves from the left (t_1) to the right side (t_6). The crack position is marked by a red dotted line. It can be seen that in t_1 , the laser spot is far away from the crack position, and the heat flux is not disrupted by the crack. The closer the spot position and crack position are, the bigger the distortion of the heat flux until the laser spot passes the crack position.

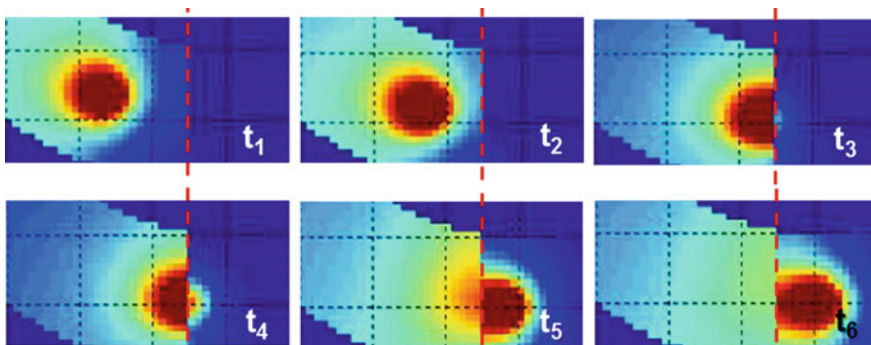


Fig. 10.13 Thermal images of a ring-shaped object with a laser spot moving across captured at different points in time. Laser moves from the left (t_1) to the right (t_6), crack position is marked by the red dashed line (colour figure online)

The approach that has been described above works well for static inspection where the camera and the object stay static until the image sequence is fully captured. If the camera or the object is moving during the acquisition phase, which is necessary in order to cover larger or complex-shaped objects, an accurate registration between the images is required. Image registration can be performed by using tracking methods. Besides sufficient overlapping areas between consecutive images, this method requires sufficient number of traceable features within every single image. Cracks or other voids can be used for feature tracking. However, this method fails in case of the lack of sufficient image features. In case of thermal imaging, this may be the case if there are no cracks or visible. A more general approach for image registration deals with the acquisition of additional time and position information for each single image. Using this information, the offset between two consecutive images can be exactly determined.

10.4.2 Process Model

The process model for a robot-based thermographic crack detection needs to deal with the heat dissipation, and therefore, material and heat source properties are the main input values. The critical question which the process model has to answer is as follows: for which areas does the crack detection work, given the position of laser, camera and test sample. Of course, this includes the question which areas are visible for the thermal camera and are at the same time in a direct light path of the laser. This covers optical aspects such as the camera's field of view and focal depth. However, additional thermal aspects such as thermal conductivity, emission coefficient, laser power, velocity, and distribution and geometry of the test sample need to be considered.

For the case of an object with flat surface and homogeneous material placed at an ideal angle under the camera and laser, the inspectable area will have a ring-shaped form as shown in Fig. 10.14. At the centre of the ring, there is the laser spot which results in a bright, white region in the image, where the pixels of the thermocamera are fully saturated and no analysis is possible. As the heat dissipates isotropically,




Fig. 10.14 Basic concept for the area that can be inspected for laser-induced heat flux evaluation (from [36])




Fig. 10.15 Projection of the process model results onto the 3D shape of the parts surface. *Left* projection on parts surface. *Right* reduction to feasible area

the signal becomes weaker until no more contrast is achieved. This area defines the outer edge of the ring-shaped region.

During the inspection process, the ring-shaped inspection area is moving across the surface leaving a trail where the part has been inspected. Cracks will become visible as a temperature gradient that is caused by interruption of the heat flux. It should be noted that the sensitivity of the detection depends on the orientation of the crack relative to the laser spot. If the crack is propagating radially from the laser spot, it will not be visible as there is no heat flow across the crack.

In the more realistic case of a non-flat part, the situation becomes significantly more complex. The model has to consider that laser and camera are not placed in an ideal position relative to the part's surfaces and that the heat propagates in a non-flat area. An approximation of the area that can be checked may be obtained by projecting the ring-shaped region onto the part's 3D surface. Additionally, self-occlusions of the part have to be considered as well as areas of high curvature, where the above-mentioned approximation is invalid. Those areas have to be excluded from the region that can be checked. Figure 10.15 shows one projection.

10.4.3 Image Analysis

Once the thermal image sequences are acquired, crack detection analysis is performed using the captured thermal images. Basically, there are two different approaches for crack analysis: the first deals with methods applied on single thermal images, and the second approach takes the temporal aspects of heat flux phenomena into account. Typically, the crack analysis is split into several processing step:

- The first step or pre-processing step deals with non-uniform correction (NUC), bad pixel replacement and calculation of temperature values if required.
- Second, the image analysis step works for single thermal images and deals with image registration, definition of area of interest (AOI), exclusion of the laser spot centre, gradient calculation and (crack) pixel classification.

- Temporal analysis step deals with the accumulation of crack pixels and morphological processing, e.g. closing algorithm, segmentation and classification.
- The final analysis step covers the transfer of resulting data into real-world coordinates and displays the results in a 3D model of the testing sample.

Above-mentioned processing steps are typically applied using intensity (“grey”) values rather than real temperature values. Since “grey” values are represented in 16-bit integer values, while temperature values are typically represented as floating point values, this approach significantly decreases processing time and memory requirements compared with the usage of temperature values.

When using infrared cameras before applying any image analysis or classification methods, prior data pre-processing is required. The pre-processing step is basically equal for all thermographic applications. Caused by the manufacturing process of thermal focal plane arrays (FPA), each pixel has its own gain and offset values. The non-uniform correction (NUC) aims at the determination of these individual gain and offset values. This can be done from two thermal images, one captured at low temperature and the second captured at high temperature. The calculated per pixel offset and gain values are then used as input values for NUC applied for all subsequent captured images. For some of the sensor pixels, typically less than 1 %, the NUC does not work. These so-called bad pixels need to be identified and replaced by an interpolation between neighbouring values. This process is called bad pixel replacement (BRP). Different kinds of thermal cameras require different intervals for renewal of NUC and BRP. For bolometer cameras, the interval for recalculation of NUC and BPR is typically some minutes, and for quantum detector cameras, e.g. with an InSb-Chip, the required NUC renewal interval is typically some hours. For some applications, the temperature values are required for subsequent analysis. Calculation of temperature values can be done by using either a parametric model or calibration data. Temperature calculation is simple for surfaces with emission coefficients of 1, which is only true for black-bodies. For objects with lower emission coefficients, like metals, the actual emission coefficient needs to be determined or measured in order to calculate proper temperature values. Once the pre-processing step has been finished, the heat flux analysis can be performed.

The image analysis step for laser-induced crack detection deals with the analysis for heat flux discontinuities or interruptions. The area with best sensitivity is described by the highest heat flux values. Figure 10.16 shows a typical temperature distribution after laser excitation (right) and the corresponding heat flux values (left). Sensitivity for crack detection is high for high heat flux values and low for areas with low heat flux values. Since the heat flux is weak for the laser spot centre and for remote areas, the crack analysis focuses on a ring-shaped area with the laser spot as its centre position.

Identification of the evaluation area starts with the detection of the laser spot position. This can be performed either by the usage of coordinate transformation methods or simply by hot spot localization. Once the spot position and its size have been identified, the surrounding area can be defined as AOI and further processed




Fig. 10.16 Local heat flux (*Left*) and temperature distribution (*Right*) after laser spot heating

for crack analysis [38]. For crack analysis, several methods that are gradient based are feasible. As described in [38] radial gradient, tangential gradient and edge-based crack detection methods are used for heat flux analysis. Besides gradient amplitude, gradient direction can be used for discrimination between cracks and signal distortions caused by effects not wanted to cause a crack signal.

The temporal heat flux evaluation takes into account that signals caused by real heat flux are strictly time dependent as described by heat conduction. Signals caused by distortions like, e.g., reflections may not show this expected temporal behaviour. Temporal evaluation therefore supports noise suppression and more accurate crack detection. Furthermore, all results gained from single image analysis step are accumulated into a joint result image. Since all AOIs are defined in image coordinates, the results need to be transformed into real-world coordinates prior to result accumulation. Since cracks are visible in more than one AOI, typically within 10 consecutive AOIs, real cracks are amplified by adding the crack signals for each single image analysis to the result image, while artefacts occurring in one single image will be diminished. This leads to an optimized dynamic range between real crack signals and artefacts caused thermal signals.

The so-calculated resulting crack image shows single crack points coming from the results of the single image analysis. In the majority of cases, these crack points can be grouped into some kind of line shapes. For some circumstances, the crack line shape may be interrupted, maybe caused by signal noise or physical properties of the crack. Cracks are of a natural form with irregularities of shape, size, width and depth. For example, for areas with fewer crack width, the resulting image may show some interruptions although the crack physically is not interrupted but show varying widths. Using a morphologic closing algorithm, such gaps can be closed, while single defect pixels are eliminated. Crack length measurement and comparison with pre-defined minimum crack length values lead to an overall pass or fail decision. Figure 10.17 illustrates a sample of gap closing for line-orientated crack pixels.

The final processing step transforms the resulting image into real-world coordinates of the test sample. This transformation results in a 3D shape of the test sample displaying location and geometry of all detected cracks.




Fig. 10.17 Crack pixels (points) with *black* connection when line segment is closed within result image and *red* connection when points are missing (colour figure online)

10.4.4 Prototype Implementation

A prototype for thermographic crack detection was set up using a M710-iC Fanuc industrial robot. The robot was selected to carry a load of 50 kg so that it is suitable for the crankshaft. A water-cooled 100 W CO₂ laser is used as heating source. With respect to the available image formats, the camera type ImageIR[®] 8300 with a pixel number of 640 × 512 was chosen. This detector size provides an adequate field of view to cover the complete test sample surface with a certain geometrical resolution in an acceptable scanning time.

Figure 10.18 shows a picture of the prototype implementation of the thermographic crack detection system for crankshafts. The prototype system runs at an




Fig. 10.18 Prototype implementation: thermographic crack detection for a crankshaft. The picture shows the robot holding the crankshaft, the thermal camera and the laser

optical resolution of 78 μm per pixel, using an effective image size of 320×256 pixels (half frame) and a field of view of 25×20 mm and capturing at a frame rate of 200 Hz. The distance between camera and crankshaft surface is 200 mm with a tolerance of ± 6 mm. Laser power used for unprocessed surfaces is between 5 and 10 W, and for processed surfaces with lower emissivity values, a laser power of 50 W is used at a scanning speed between 60 and 200 mm/s. For safety reasons, the workcell is made of hard metal sheets and a security switch which allows laser operation only when doors are properly closed.

10.5 Conclusions and Outlook

From the first implementations of industrial inspection robots for quality control, much progress has been made in terms of increasing the level of autonomy and making such complex systems easier to handle. The solution of complex inspection problems is already feasible today, but still requires significant effort for the implementation of a particular application. Semi-automatic methods for setting up the inspection process help to reduce this effort, but have not yet reached a level of abstraction that would allow generic solutions for whole classes of applications. The structures and examples presented in this chapter provide a basis for such generic solutions that require specific modules, such as augmented CAD data of the part, a process model for the inspection process and a 3D model of the workcell to set up a whole inspection task. Future research will enable an easy exchange of the single modules to quickly adapt an inspection system to a new task or a new inspection process and it will not take too long, before such systems are installed in industrial production lines.

Acknowledgments The work presented in this chapter received cofunding from the European Commission in the 7th Framework Programme, projects “ThermoBot” (No. 284607) and FibreMap (No. 608768), and the Austrian Research Funding Agency (FFG), projects “SelTec”, “ProFit” and “LISP”.

The authors would like to thank all partners of these projects, and especially, Prof. Emanuele Menegatti, Dr. Stefano Ghidoni and the whole team of the IAS Lab of the University of Padova.

References

1. Tătar O, Măndru D, Ardelean I (2007) Development of mobile minirobots for in pipe inspection tasks. *MECHANIKA*, 6(68):1392–1207
2. Granosik G, Borenstein J, Hansen MG (2006) Serpentine robots for industrial inspection and surveillance. In: Huat LK (ed) *Industrial robotics: programming, simulation and application*, pp 702, ISBN: 3 86611-286-6
3. Edinbarough I, Balderas R, Bose S (2005) A vision and robot based on-line inspection monitoring system for electronic manufacturing. *Comput Ind* 56:986–996
4. Woern H, Laengle T, Gauss M (2003) ARIKT: adaptive robot based visual inspection. *Künstliche Intelligenz* 2:33–35

5. Kuhlenkoetter B, Krewet C, Schueppstuhl T (2006) Adaptive robot based reworking system, industrial robotics: programming, simulation and applications. Low Kin Huat (ed), ISBN: 3-86611-286-6
6. Yang CC, Ciarallo FW (2001) Optimized sensor placement for active visual inspection. *J Robotic Syst* 18(1):1–15
7. Biegelbauer G, Vincze M, Noehmayer H, Eberst C (2004) Sensor based robotics for fully automated inspection of bores at low volume high variant parts. *IEEE international conference on robotics and automation*, 5:4852–4857, 26 April–1 May 2004
8. Merat FL, Radack GM (1992) Automatic inspection planning within a feature-based CAD system. *Robotics Comput Integr Manuf* 9(1):61–66
9. Park TH, Kim HJ, Kim N (2006) Path planning of automated optical inspection machines for PCB assembly systems. *Int J Control Autom Syst* 4(1):96–104
10. Scoot WR, Roth G (2003) View planning for automated three-dimensional object reconstruction and inspection. *ACM Comput Surv* 35(1):64–96
11. Lee KH, Park HP (2000) Automated inspection planning of free-form shape parts by laser scanning. *Robot Comput Integr Manuf* 16(4):201–210
12. Lu CG, Morton D, Wu MH, Myler P (1999) Genetic algorithm modelling and solution of inspection path planning on a coordinate measuring machine (CMM). *Int J Adv Manuf Technol* 15:409–416
13. Scott W, Roth G, Rivest JF (2002) Pose error effects on range sensing. In *Proceedings of the 15th international conference on vision interface (Calgary, Alta., Canada)*, pp 331–338
14. Pisinger D, Ropke S (2010) Large neighborhood search. *Handbook of Metaheuristics*, pp 399–419
15. Dorigo M, Maniezzo V, Colomi A (1996) The ant system: optimization by a colony of cooperating agents. *IEEE Trans Syst Man Cybern Part B* 26(1):29–41
16. Ankerl M, Hämmerle A (2009) Applying ant colony optimisation to dynamic pickup and delivery. In: *Comput Aided Syst Theory-EUROCAST*, pp 721–728
17. Daniilidis K (1999) Hand-eye calibration using dual quaternions. *Int J Robot Res* 18(3):286–298
18. Zhao Z, Liu Y (2006) Hand-eye calibration based on screw motions. *18th International conference on pattern recognition (ICPR'06)*, 3:1022–1026
19. Hollerbach JM, Wampler CW (1996) The calibration index and taxonomy for robot kinematic calibration methods. *Int J Robot Res* 15(6):573–591
20. Pradeep V, Konolige K, Berger E (2010) Calibrating a multi-arm multi-sensor robot: a bundle adjustment approach. In: *Proceedings of the International symposium on experimental robotics (ISER)*, Delhi India, Dec 18–21 2010
21. Strobl KH, Hirzinger G (2008) More accurate camera and hand-eye calibrations with unknown grid pattern dimensions. *Proceedings-IEEE international conference on robotics and automation*, pp 1398–1405
22. Eitzinger C, Heidl W, Lughofer E, Raiser S, Smith JE, Tahir MA, Sannen D, Van Brussel H (2009) Assessment of the influence of adaptive components in trainable surface inspection systems. *Mach Vis Appl J*, doi: [10.1007/s00138-009-0211-1](https://doi.org/10.1007/s00138-009-0211-1)
23. Smith JE, Tahir MA, Caleb-Solly P, Lughofer E, Eitzinger C, Sannen D, Nuttin M (2009) Human-machine interaction issues in quality control based on on-line image classification. *IEEE Trans Syst Man Cybern* 39(5):960–971
24. Eitzinger C, Thumfart S (2012) Optimizing feature calculation in adaptive machine vision systems. In: *Sayed-Mouchaweh M, Lughofer E, Learning in non-stationary environments: methods and applications*, Springer Science+Business Media New York, doi: [10.1007/978-1-4419-8020-5_13](https://doi.org/10.1007/978-1-4419-8020-5_13)
25. Szeliski R (2006) Image alignment and stitching: a tutorial. *Found Trends Comput Graph Vis* 2(1):1–104
26. Brown M, Lowe D (2007) Automatic panoramic image stitching using invariant features. *Int J Comput Vision* 74(1):59–73

27. Kopf C, Heindl C, Rooker M, Bauer H, Pichler A (2013) A portable, low-cost 3D body scanning system. 4th International conference and exhibition on 3D body scanning technologies, CA USA Nov 19–20 2013
28. Schmitt R, Mersmann C, Schoenberg A (2009) Machine vision industrialising the textile-based FRP production. In: Proceedings of 6th international symposium on image and signal processing and analysis, pp 260–264
29. Palfinger W, Thumfart S, Eitzinger C (2011) Photometric stereo on carbon fibre surfaces. Proceeding of the Austrian association for pattern recognition
30. Thumfart S, Palfinger W, Stöger M, Eitzinger C (2013) Accurate fibre orientation measurement for carbon fibre surfaces. 15th International conference on computer analysis of images and patterns, pp. 75–82
31. Woodham R (1989) Photometric method for determining surface orientation from multiple images. *Opt Eng* 19(1):139–144
32. Johnson M, Adelson E (2011) Microgeometry capture using an elastomeric sensor. *Comput Vis Pattern Recogn* pp 2553–2560
33. Zhou K, Wang L, Tong Y, Desbrun M, Guo B, Shum HY (2005) Texture montage: seamless texturing of arbitrary surfaces from multiple images. Proceedings of ACM SIGGRAPH, pp 1148–1155
34. Avdelidis N, Gan T-H, Ibarra-Castanedo C, Maldaque X (2011) Infrared thermography as a non-destructive tool for materials characterisation and assessment. Proceedings—SPIE the international society for optical engineering, (8013–8039) Thermal Infrared Applications XXXIII
35. Holst G (2000) Common sense approach to thermal imaging. SPIE Volume PM-86, pp 60, ISBN: 0-8194-3722-0
36. Taib S, Jadin M, Kabir S (2012) Thermal Imaging for enhancing inspection reliability: detection and characterization, infrared thermography, Dr. Raghu V Prakash (Ed.), ISBN: 978-953-51-0242-7
37. Traxler G, Thanner P (2011) Automatisierte Wärmeflussprüfungen in der Stahlindustrie, Leitfaden zur Wärmeflussthermografie, ISBN 978-8396-0234-8
38. Ghidoni S, Minella M, Nanni L, Ferrari C, Moro M, Pagello E, Menegatti E (2013) Automatic crack detection in thermal images for metal parts. International conference on heating by electromagnetic sources (HES-13)