Advancement in Neural Radiance Fields Propulsion for SLAM: A Novel Upgrade

Ruotong Sun [1*], Jiachen Lu[2] and Yankun Chen [3]

[1] Mechanical Engineering Institute, Xi'an Jiaotong University, Xi'an,710049, China
[2] CQU-UC Joint Co-op Institute, Chongqing University, Chongqing,400044, China
[3] Artificial Intelligence Institute, Jianghan University, Wuhan,430056, China
srt040416@stu.xjtu.edu.cn

Abstract. In recent years, simultaneous localization and mapping (SLAM) techniques have seen continual improvement. Meanwhile, the development of Neural Radiance Fields (NeRF) has surged, demonstrating significant advantages in the mapping process at the backend of SLAM. Researchers are increasingly exploring new approaches integrating NeRF deeply into SLAM technology, propelling SLAM towards greater intelligence and efficiency. Compared to traditional SLAM methods, NeRF-based SLAM offers numerous advantages, albeit facing several challenges. This paper reviews the development of NeRF technology and analyzes representative NeRF-based SLAM methods. Through comparison and analysis, it summarizes the strengths and weaknesses of NeRF-based SLAM methods, and anticipates the future directions and application prospects of NeRF technology in the SLAM field, highlighting research directions such as real-time improvement, performance optimization, and interdisciplinary integration. This paper aims to provide researchers with insights into NeRF-based SLAM technology, fostering further research and application in this field.

Keywords: SLAM, NeRF, Robotics.

1 Introduction

With the rapid development of technology, SLAM has become an indispensable core technology in fields such as robotics, autonomous driving, and augmented reality [1]. The SLAM system aims to achieve autonomous navigation of robots in unknown environments by simultaneously constructing environment maps in real-time and determining their own positions, providing precise navigation information for robots. This is vital for advancing sectors including autonomous driving, smart homes, and various medical procedures, both non-invasive and minimally invasive

A classic SLAM typically consists of five stages: mapping, tracking, optimization, loop closure, and localization. Traditional SLAM methods achieve localization and map construction based on geometric principles and feature point matching by extracting features from sensor data. However, in environments with low texture or significant illumination changes, where there are few feature points available for

© The Author(s) 2024
Y. Wang (ed.), *Proceedings of the 2024 International Conference on Artificial Intelligence and Communication (ICAIC 2024)*, Advances in Intelligent Systems Research 185,
https://doi.org/10.2991/978-94-6463-512-6_16

extraction and matching, the performance of traditional SLAM may be suboptimal [1]. Given the demand to overcome current challenges and the significant achievements of neural network technology in computer vision in recent years, especially in image generation and recognition, the fusion of deep learning and traditional geometric methods has become a new trend in SLAM. Considering one of the key challenges faced by SLAM is the ability to synthesize and reinforce real-time retrieval and application of environmental information from new perspectives, the development of NeRF in recent years has led to the integration of NeRF and SLAM becoming a frontier development hotspot in the SLAM field [2]. Unlike traditional geometry-based rendering methods, NeRF directly manipulates raw pixel values to generate photo-realistic maps without relying on pre-extracted features or geometric information. This makes it possible to directly perform localization and map construction from image data using neural networks. The application of NeRF in the SLAM field shows promising prospects, while the research process has also revealed certain limitations.

To address the inadequate attention to the latest developments in SLAM surveys and meet the growing interest in research in this field, this paper reviews the past four years of NeRF-based SLAM technology. In-depth analysis of representative NeRF-integrated SLAM systems such as iMAP, Nice-SLAM, NoPe-NeRF etc, presents the rapid progress in this field [3][4][5]. In addition, unlike mere enumeration and comparison, this review introduces these studies from a novel perspective by utilizing the integration points of NeRF and SLAM frameworks as the classification basis. It provides a comprehensive overview of the incorporation of NeRF in the five key stages of mapping, tracking, optimization, loop closure, and localization within SLAM. This encompasses an analysis of representative research at each stage along with their underlying principles, applications, advantages, and challenges. Through thorough investigation and comparison of these methodologies, this paper elucidates the strengths and limitations associated with NeRF-based SLAM while also discussing potential future directions for development.

2 The Overview of NeRF

2.1 Mapping

SLAM requires maps that are continuously expanding and easy to query. On the other hand, the process of building maps with NeRF is easily scalable, convergent, and superior in rendering speed. The representation of NeRF maps can be categorized into three types: implicit, explicit, and hybrid. Implicit refers to the use of Multilayer Perceptrons (MLPs) to store the map, while explicit means using traditional methods such as voxel grids to represent the map, and hybrid representation refers to a combination of implicit and explicit methods. Implicit maps can be further divided into maps represented by a single Multilayer Perceptron (MLP) and maps represented by multiple MLPs. Below, this paper will analyze each in detail.

Implicit Representation. The iMAP, published in 2021, is a pioneering work in the field of NeRF-based SLAM, which uses a single MLP for map representation [3]. To meet the needs of real-time mapping, iMAP has only four hidden layers, each with 256 neurons, and does not consider reflections, so it does not input the viewing direction. Due to its single MLP, it suffers from a severe forgetting problem. However, in this system, researchers addressed the forgetting problem by organizing keyframes into a memory bank set for continuous backend map optimization, to some extent. From the map effect, it can be seen that the single MLP's representation capability is limited and can only model small scenes, leading to new challenges for applying NeRF to SLAM systems—modeling unbounded scenes.

Mip-NeRF 360 is the first attempt of NeRF to model unbounded scenes [6]. Its predecessor, Mip-NeRF, was used to improve rendering quality and speed up rendering, while Mip-NeRF 360 extended the Kalman filter to map unbounded scenes into bounded coordinate spaces, and then mapped the set of mappings to model unbounded scenes.

NeRF++ is another attempt of NeRF on unbounded scenes [7]. It divides the space into foreground and background, reconstructs the foreground using the original NeRF method, and then uses the method of reverse spherical parameterization to build the background on an outer sphere. However, this solution and the unbounded scenes processed by Mip-NeRF 360 are around a central camera running a week to establish a model, while the camera pose of SLAM is arbitrary. To achieve unbounded scene modeling under arbitrary camera poses, the multiple MLP implicit map representation emerged.

Block-NeRF is a representative model based on multiple MLPs for implicit map representation [8]. It divides the scene into multiple separate blocks, which are set at key positions throughout the scene to ensure that the visible range of these blocks can cover the entire scene. These blocks are trained separately, and each block can be updated individually, avoiding updating the entire map each time. During rendering, the visible blocks within the visible range are superimposed according to visibility to complete map rendering. However, distant blocks do not participate in viewpoint synthesis, so this model does not truly solve the problem of background modeling, but its strategy is close to the needs of SLAM local maps.

Another model of multiple MLP implicit map representation is Mega-NeRF, which achieves scalable reconstruction of large scenes [9]. This model divides the entire scene into regular spatial cells, each with its MLP weight. At the same time, it budgets the spatial cell density and RGB based on the idea of Plenoctress, placing them in an octree to improve rendering speed [10]. However, it upgrades Plenoctress's fixed octree to a scalable octree to save computational resources. In addition, for the handling of foreground and background, Mega-NeRF refers to the strategy of NeRF++. While NeRF++ uses a sphere to encapsulate the camera and objects in the foreground, resulting in many invalid sampling points on the ground, Mega-NeRF changes the sphere of the foreground to an ellipsoid, reducing such invalid sampling and thus reducing computational consumption. However, even so, the modeling speed of the NeRF system still cannot meet the real-time needs of the existing SLAM.

In summary, if the backend map of the SLAM system adopts an implicit scene representation, the currently available mapping strategies mainly include accelerating training, accelerating rendering, and map expansion. For accelerating the training process, if it is based on a single MLP map representation, acceleration methods such as Instant-NGP can be used to improve the performance of MLP itself; if it is based on multiple MLP map representations, some MLPs can be selected for training according to certain strategies and updated locally, such as Block-NeRF [11][8]. Accelerating the rendering process can be achieved by increasing the proportion of effective sampling points, or by selecting some MLPs for rendering according to certain strategies, which is faster than global rendering. In addition, pre-extracting the body density and RGB and building an octree as a buffer like Mega-NeRF or plenoctree can avoid querying the MLP during rendering, thereby improving rendering speed [9][10]. For map expansion, the current solution is to pre-divide the area and call multiple NeRFs autonomously allocated by the SLAM system based on the camera's running status. However, this map expansion method is limited by significant pre-planning and cannot achieve dynamic expansion.

Therefore, the current implicit scene representation method of NeRF can significantly improve the rendering quality of scene reconstruction, but the problems of slow speed and difficulty in expansion still need to be solved. It still cannot fully meet the needs of SLAM backend mapping.

Explicit Representation. Explicit scene representation does not rely on MLPs, but since NeRF's application to SLAM is crucial to achieve volume rendering, as long as the map supports volume rendering, NeRF's application to SLAM can be realized.

DVGO is a typical model of explicit representation, which, detached from MLPs, directly implements voxel grid optimization, achieving fast convergence of radiance field reconstruction [12]. Compared to the original NeRF, its training speed has increased by two orders of magnitude. When rendering, DVGO first extracts the voxel density values around the specified position, then performs trilinear interpolation to obtain the density value at position x, activates it through softplus, and calculates the opacity through the alpha formula. Although this type of map representation optimization has the advantages of fast convergence speed, easy expansion, and editing, it is difficult to fit high-frequency features, and the operation of volume rendering is challenging, so researchers have begun to explore mixed scene representation methods of explicit and implicit.

Hybrid Representation. Nice-SLAM truly achieves mapping with a mixed representation of explicit and implicit scenes, inheriting the camera tracking idea of iMAP, but it remedies the shortcomings of iMAP's failure in large-scale scenes [4][3]. Nice-SLAM uses three voxel grids of different resolutions, nested to represent the scene, storing feature vectors in the voxel grid, sampling with trilinear interpolation, and then decoding with a pre-trained decoder to merge the decoding content of each layer of the grid to achieve volume rendering. However, since its voxels are pre-allocated, it cannot achieve map expansion.

Vox-fusion is another NeRF-based mixed scene representation SLAM method, which realizes map expansion and reduces the memory consumption of Nice-SLAM from 200M to 0.15M [13][4]. Its map is stored in an octree, which can expand continuously with the progress of SLAM. The grid stores the encoding of the signed distance field (SDF). Vox-fusion sets the SDF value of the object surface position to 0 in the grid. The surface is positive outward and negative inward, thus achieving excellent fitting of the scene's geometric structure. However, both Nice-SLAM and Vox-fusion lack new viewpoint synthesis capabilities in map representation, resulting in gaps in map representation. In addition to voxels, the explicit scene representation of mixed schemes can also return to a more traditional point cloud-based scene representation, which stores feature vectors in the point cloud and achieve continuity through interpolation.

The most intuitive role of NeRF in SLAM is as the backend scene representation. In summary, NeRF-based SLAM map representations can be divided into three categories: using one or multiple MLPs to store the full implicit representation of the scene. Its advantages are that it can continue the excellent characteristics of the original NeRF, with continuous scene representation, convenient optimization, but in practice, they do not quite meet the needs of SLAM for the backend because such maps are often difficult to expand and cannot be applied in real-time. The second type is a fully explicit voxel grid, which is less frequently used due to its difficulty in fitting high-frequency features and not being very friendly to volume rendering. However, its advantages, fast optimization convergence speed, easy expansion, and editing, have inspired NeRF work. The third type is based on the explicit storage of voxel or point cloud feature vectors, which are then decoded by MLPs during rendering. The advantage of this approach is that it retains the advantages of traditional map representation—easy expansion, editing, training, and fast rendering speed—and can achieve rendering quality to some extent. In conclusion, for mapping, the mixed scene representation based on voxels and feature vectors can meet the needs of SLAM and is the most suitable and commonly used.

2.2 Tracking

In SLAM, precise, fast, and robust camera pose estimation is required at the frontend. The NeRF map, acting as the backend, leverages its inverse rendering capability to aid in pose estimation at the frontend. Nice-SLAM achieves camera pose optimization and local map updates simultaneously through NeRF's inverse rendering [4]. Furthermore, Nicer-SLAM, as an improvement, surpasses Nice-SLAM in terms of localization accuracy and rendering quality [14][4]. However, NeRF inverse rendering has its drawbacks, particularly sensitivity to initial pose. When the initial pose deviates significantly, the accuracy of local map optimization achieved by NeRF inverse rendering is greatly reduced. Subsequently, NoPe-NeRF proposes a method to enhance camera pose estimation accuracy during training by introducing undistorted monocular depth priors [5]. These priors, generated to correct scale and translation parameters, constrain the relative poses between adjacent frames. To enforce these constraints, the research team introduces novel loss functions, including Chamfer distance loss based on undistorted depth maps generated from monocular depth and depth-based surface rendering loss. Experimental results demonstrate that this method performs excellently in handling challenging camera trajectories and outperforms earlier methods in terms of new viewpoint rendering quality and pose estimation accuracy.

2.3 Optimization

In NeRF, camera pose optimization is also necessary, as inaccurate camera poses can pose challenges to reconstruction. Synchronous reconstruction and optimization are mutually beneficial. BARF introduces image alignment theory into NeRF, jointly optimizing camera poses and reconstruction models [15]. Therefore, when adjusting camera poses, it uses low-pass filters at different frequencies, equivalent to aligning low-frequency image information for pose optimization.

Additionally, VBA implements optimization on NeRF [16]. It performs online scene reconstruction, employing a B+ tree combined with VDB grids to store maps. This can be seen as an upgraded version of octrees combined with regular voxel grids, enabling more efficient expansion and querying. Its grids store multi-layer spatial features, decoded through MLP, and the optimization method is similar to the inverse rendering approach. It simultaneously optimizes camera poses, features stored in the grid, and MLP parameters by calculating the loss between the rendered RGB, Depth, and ground truth.

2.4 Loop Closure

In SLAM systems, accumulated drift errors from sensor data collection and long-term camera pose estimation necessitate loop closure detection modules to correct drift errors by detecting whether the camera has passed through previously visited locations. A classic SLAM system ORB-SLAM2 utilizes a bag-of-words model specifically for matching ORB features of the current frame to identify loop closures. After identifying loop closures, correction of keyframe poses is necessary. However, for NeRF, the backend map encodes keyframe image content within the map, especially for implicit maps. This implies that the mapping process is irreversible, and the encoded information cannot be adjusted for pose correction, thus loop closure correction cannot be achieved.

SE3 addresses this issue by implementing an equivariant mapping from SE3 to feature space, synchronously adjusting camera poses and implicit map representations [17]. Leveraging the frontend localization module of ORB-SLAM2, the authors utilize their voxel grid as the backend, constructing a SLAM system that achieves loop closure correction for implicit maps.

2.5 Localization

The foundation of loop closure detection is achieving camera localization on the global map. On top of this, relocalization can restart SLAM after camera tracking loss or interruption.

IR-MCL achieves global localization in NeRF maps built from 2D radar scans through Monte Carlo localization [18]. It samples several candidate camera poses in space, calculates their differences with observed values, generates weights, eliminates candidate poses with low weights, iteratively updates poses, and finally, when converging near the true pose, weights the remaining poses to obtain the final pose. This idea may be applicable to 3D maps, but whether convergence is achievable in 3D space requires further investigation.

CP-SLAM proposes a novel neural network-based 3D scene representation method within NeRF, where each point maintains a learnable neural feature for scene encoding, associated with specific keyframes [19]. Additionally, it updates the global optimization framework, similar to traditional bundle adjustment, improving system accuracy. This enhances both localization and reconstruction capabilities compared to existing methods.

3 Discussion

3.1 Integration of Traditional Point Clouds and Neural Radiance Fields

The point cloud-based mapping approach combines traditional point clouds with neural radiance fields by storing feature vectors in point clouds. This integration not only achieves NeRF rendering quality but also retains the intuitive and easily scalable characteristics of point clouds. Importantly, once the map returns to a point cloud, various backend optimization methods of traditional SLAM can be utilized, which is significant for SLAM. However, there remains a gap in reapplying traditional SLAM's backend optimization methods to NeRF-based SLAM research.

3.2 Migration Challenges of Geometry-Based Optimization Algorithms

Compared to traditional geometry-based map representations, NeRF faces challenges in achieving frame-to-map alignment. Traditional maps are typically based on geometric features such as keypoints or edges, easily aligned with current frames using traditional geometric matching algorithms (e.g., feature matching, ICP, etc.). However, since NeRF maps are dense voxel representations learned by neural networks rather than sparse feature points or geometric structures, traditional geometry-based matching methods cannot be directly applied.

Therefore, NeRF map pose optimization often employs iterative methods to minimize reprojection errors between current frames and the map. This optimization process may require additional information, such as optical flow or depth information, to aid in pose estimation. However, due to the unique representation of NeRF maps, optimizing them is not as straightforward as traditional maps, especially in scenes with rich features and textures, which may lead to problems such as local optima and slow convergence. Therefore, migrating mature geometry-based optimization algorithms to NeRF maps is a challenging task.

3.3 Adaptive NeRF-based SLAM Methods for Dynamic Environments

In adaptive NeRF-based SLAM methods for dynamic environments, we face various challenges. The presence of dynamic objects may interfere with traditional NeRF methods, leading to instability in reconstruction results and increased localization errors. Therefore, to enhance adaptability in dynamic environments, modeling and tracking mechanisms for dynamic objects can be introduced. By modeling dynamic objects and tracking them with motion models, they can be distinguished from static scenes and dynamically updated in the NeRF reconstruction process to maintain reconstruction accuracy. This approach effectively addresses challenges posed by dynamic environments, making NeRF-based SLAM methods more robust and reliable.

Furthermore, objects in dynamic environments often exhibit certain temporal characteristics, with their motion trajectories and appearances changing over time. Therefore, leveraging temporal information can better understand and reconstruct dynamic environments. Combining visual SLAM techniques allows for the inference of dynamic object motion trajectories between consecutive frames and consideration during NeRF reconstruction. By fully utilizing temporal information, the stability and accuracy of NeRF-based SLAM in dynamic environments can be improved, providing more reliable solutions for practical applications.

3.4 Trend of Multi-Sensor Fusion

Under the trend of multi-sensor fusion, incorporating multi-view information processing and multi-sensor fusion can improve system performance. In dynamic environments, single-view image data may be affected by dynamic object occlusion and interference, impacting reconstruction results. To address this issue, multi-view information can be introduced for processing. By fusing image data from multiple views, a more comprehensive environmental information capture can be achieved, reducing reconstruction errors caused by occlusion and interference, thus improvingrobustness and accuracy of reconstruction. This multi-view information processing strategy effectively addresses challenges in dynamic environments, providing more reliable support for NeRF-based SLAM methods. Additionally, multi-sensor fusion technology combines different sensor data, such as depth cameras, IMUs, LiDAR, etc., to obtain richer environmental information and enhance reconstruction and localization accuracy and robustness through well-designed fusion strategies. Therefore, multi-sensor fusion technology is of great significance for improving the performance of NeRF-based SLAM in dynamic environments and deserves further research and exploration.

4 Conclusion

With the continuous development of neural radiance field methods (NeRF) and other neural rendering-based environment reconstruction methods, the integration of NeRF models into SLAM systems has significantly advanced the technology, notably enhancing continuous mapping, robustness, and generalization capabilities, as well as addressing many shortcomings of traditional SLAM models. NeRF-based SLAM methods have shown tremendous potential in environmental reconstruction and localization. This paper comprehensively analyzes over a dozen NeRF-based SLAM papers in the order of various stages of SLAM technology, studying the evolution and advantages and disadvantages of their methods, discussing the improvement of system performance by running dual-thread parallelism in NeRF-based SLAM systems, and discovering some highly mature frameworks. Future research will largely build upon these established foundations. Finally, the paper points out the problems existing in current NeRF-based SLAM systems, including real-time performance, adaptability to dynamic environments, and multi-sensor fusion, which still need further research and solutions. Through this summary, the author hopes to attract more attention from relevant practitioners, promote research on NeRF-based SLAM, and further advance the development of future SLAM systems.

Authors Contribution

All the authors contributed equally.

References

1. Cheng, J., Zhang, L., Chen, Q., Hu, X., & Cai, J.: A review of visual SLAM methods for autonomous driving vehicles. Engineering Applications of Artificial Intelligence **114**, 104992 (2022)
2. Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. Communications of the ACM **65**(1), 99-106 (2021).
3. Sucar, E., Liu, S., Ortiz, J., Davison, A. J.: iMAP: Implicit mapping and positioning in real-time. In Proceedings of the IEEE/CVF International Conference on Computer Vision 6229-6238 (2021).
4. Zhu, Z., Peng, S., Larsson, V., Xu, W., Bao, H., Cui, Z. Pollefeys, M.: Nice-slam: Neural implicit scalable encoding for slam. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition 12786-12796 (2022).
5. Bian, W., Wang, Z., Li, K., Bian, J. W., Prisacariu, V. A.: Nope-nerf: Optimising neural radiance field with no pose prior. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2023).
6. Barron, J. T., Mildenhall, B., Verbin, D., Srinivasan, P. P., Hedman, P.: Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields Supplemental Materials (2023).
7. Zhang, K., Riegler, G., Snavely, N., Koltun, V.: Nerf++: Analyzing and improving neural radiance fields. arXiv preprint arXiv (2020).
8. Tancik, M., Casser, V., Yan, X., Pradhan, S., Mildenhall, B., Srinivasan, P. P., Kretzschmar, H.: Block-nerf: Scalable large scene neural view synthesis. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022).
9. Turki, H., Ramanan, D., Satyanarayanan, M. Mega-nerf: Scalable construction of large-scale nerfs for virtual fly-throughs. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition pp. 12922-12931 (2022)

10. Yu, A., Li, R., Tancik, M., Li, H., Ng, R.,Kanazawa, A.: Plenoctrees for real-time rendering of neural radiance fields. In Proceedings of the IEEE/CVF International Conference on Computer Vision (2021).
11. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM transactions on graphics (TOG) **41**(4), 1-15 (2022).
12. Sun, C., Sun, M., & Chen, H. T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2022).
13. Yang, X., Li, H., Zhai, H., Ming, Y., Liu, Y., & Zhang, G. (October). Vox-fusion: Dense tracking and mapping with voxel-based neural implicit representation. In 2022 IEEE International Symposium on Mixed and Augmented Reality (ISMAR) (2022).
14. Zhu, Z., Peng, S., Larsson, V., Cui, Z., Oswald, M. R., Geiger, A.,Pollefeys, M.: Nicer-slam: Neural implicit scene encoding for rgb slam. (2023).
15. Hedman, P., Srinivasan, P. P., Mildenhall, B., Barron, J. T., & Debevec, P.: Baking neural radiance fields for real-time view synthesis. In Proceedings of the IEEE/CVF International Conference on Computer Vision pp. 5875-5884. (2021).
16. Clark, R. Volumetric bundle adjustment for online photorealistic scene capture. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition pp. 6124-6132. (2022).
17. Yuan, Y., Nüchter, A.: An algorithm for the SE (3)-transformation on neural implicit maps for remapping functions. IEEE Robotics and Automation Letters **7**(3), 7763-7770 (2022).
18. Kuang, H., Chen, X., Guadagnino, T., Zimmerman, N., Behley, J., Stachniss, C.: Ir-mcl: Implicit representation-based online global localization. IEEE Robotics and Automation Letters **8**(3), 1627-1634 (2023).
19. Hu, J., Mao, M., Bao, H., Zhang, G., Cui, Z. Cp-slam: Collaborative neural point-based slam system. Advances in Neural Information Processing Systems, **36** (2024).

Open Access This chapter is licensed under the terms of the Creative Commons Attribution-NonCommercial 4.0 International License (http://creativecommons.org/licenses/by-nc/4.0/), which permits any noncommercial use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.