

# Beyond Static Gaussians: An Empirical Investigation of Architectural Paradigms for Dynamic 3D Scene Reconstruction

Adrian Ramlal  
University of Waterloo  
adrian.ramlal@uwaterloo.ca

John S. Zelek  
University of Waterloo  
jzelek@uwaterloo.ca

## Abstract

*Dynamic scene reconstruction via 3D Gaussian Splatting (3DGS) has emerged as a compelling approach for representing evolving environments, yet understanding trade-offs between methodologies remains crucial. This paper presents a comprehensive analysis of dynamic 3DGS methods, categorizing them into two paradigms: structure-guided methods employing auxiliary representations (deformation fields, canonical spaces, grids) to model temporal changes, and gaussian-centric methods encoding dynamics directly into primitives via continuous functions or 4D representations. We evaluate representative methods from both paradigms on the D-NeRF benchmark. Our findings reveal that structure-guided methods achieve superior reconstruction fidelity and compact model sizes, while gaussian-centric approaches demonstrate significantly higher rendering speeds enabling real-time performance, though with greater quality variability and potentially substantial storage overhead. This analysis highlights a fundamental trade-off between reconstruction quality/compactness versus rendering speed, providing insights to guide future research and application development in dynamic scene reconstruction.*

## 1. Introduction

Dynamic scene reconstruction captures and digitally represents environments where geometry and appearance evolve over time [18], underpinning applications in VR/AR, interactive media, robotics, and autonomous navigation. However, the task is inherently sparse and under-constrained, relying on limited viewpoints and time samples that leave gaps in data [2]. The complexity of modeling non-rigid deformations, abrupt changes, and lighting variations further complicates reconstruction.

3D Gaussian Splatting (3DGS) [11] has emerged as a powerful solution, offering explicit scene representation and efficient real-time rendering. Methods for adapt-

ing 3DGS to dynamic scenes have coalesced into two paradigms [3]: *structure-guided* methods leverage auxiliary structures, MLP-based deformation fields [24] or sparse control points [26], to explicitly model temporal transformations; *gaussian-centric* methods encode dynamics directly into primitives via continuous functions [10, 15], Radial Basis Functions [13], or 4D space-time representations [5, 26].

Our comparative analysis reveals distinct performance characteristics: structure-guided methods achieve higher reconstruction fidelity and temporal consistency with compact models, particularly on monocular benchmarks like D-NeRF [19], while gaussian-centric methods demonstrate superior computational efficiency and rendering speed essential for real-time interaction, though with increased storage demands. This paper evaluates these paradigms' relative merits and trade-offs in reconstruction quality, rendering speed, and model size, illuminating the current landscape of dynamic 3DGS and addressing key challenges in dynamic scene reconstruction.

## 2. Related Works

The field of dynamic 3DGS has evolved rapidly, with researchers developing diverse approaches to address the challenges of reconstructing scenes that change over time. In this section, we review the current state of the art, categorizing methods into two distinct paradigms: structure-guided methods and gaussian-centric methods.

### 2.1. Structured-Guided

Structure-guided methods leverage additional structured information such as neural networks, canonical spaces, or grid-based frameworks to model deformations in dynamic scenes. These approaches separate the representation of the scene's structure from its temporal evolution, using auxiliary models to predict how Gaussian primitives should transform temporally.

Deformable 3DGS [24] pioneered the use of deformable 3D Gaussians for dynamic scenes by learning them in a

canonical space and employing an implicit deformation field implemented as an MLP to model spatial-temporal deformation. This approach allows for complex deformations while maintaining a consistent scene structure in the canonical space. To mitigate overfitting caused by inaccurate poses, they introduce an annealing smooth training paradigm with linearly decaying Gaussian noise. GauFRE [14] extends this by decoupling dynamic and static scene modeling, applying similar deformation techniques only to the dynamic elements.

Also building on the foundation of [24], GaGS [16] introduced a method that voxelizes Gaussian distributions and utilizes sparse convolutions to extract geometry-aware features for deformation learning, enhancing the system’s ability to capture fine-grained structural changes. Another approach, DynMF [12], models dynamics by factorizing scene motion into a compact set of learnable basis trajectories, arguing that complex movements can be efficiently represented by combining these bases. For a different take on efficient representation, SC-GS [9] employs a sparse set of control points—significantly fewer than the Gaussians—to drive deformation. These control points utilize radial basis functions and Linear Blend Skinning to interpolate transformations across the dense Gaussian field, and the system adaptively prunes or creates control points based on learned parameters.

4D-GS [22] employs multi-scale HexPlane [2] as its foundational representation, encoding both temporal and spatial information through grid-based structures. This approach uses multi-head decoders to predict different Gaussians’ attributes separately, enhancing the model’s expressiveness. DeformGS [6] incorporates local rigidity and isometric loss constraints while designing regularization based on conservation of momentum to achieve smoother motion trajectories.

Structure-guided methods often benefit from incorporating various priors to address the under-constrained nature of dynamic reconstruction. Park et al. [20] demonstrate this by consolidating supervisory signals from depth maps and long-range 2D tracks to achieve consistent dynamic representations. Similarly, Guo et al. [7] develop a flow-augmentation method that leverages correspondence between 3D Gaussian motion and pixel-level flow to improve reconstruction quality.

## 2.2. Gaussian-Centric

Gaussian-centric methods encode scene changes directly into the Gaussian primitives, often without relying on auxiliary structured representations. These methods adjust Gaussian parameters over time using continuous functions or frame-by-frame updates.

Dynamic 3DGS [17] introduced one of the earliest approaches in this category, where Gaussians maintain con-

sistent appearance properties (color, opacity, size) while their spatial properties (position, orientation) evolve temporally. Their online reconstruction process uses each frame to initialize the next, applying physical constraints like local rigidity, local rotational-similarity, and long-term local-isometry to ensure plausible Gaussian motion and rotation.

Continuous temporal modeling emerged as a key advancement: Katsumata et al. [10] incorporated Fourier approximations over time and optical flow supervision to model attribute changes of 3DGS, ensuring continuous changes without introducing excessive parameters. ED3DGS [1] introduced per-Gaussian latent embeddings alongside temporal embeddings.

This continuous representation trend continued with STG [13], which utilized Radial Basis Functions for modeling opacity changes and parametric polynomials for position and rotation. Gaussian-Flow [15] combined time-domain polynomials with frequency-domain Fourier series to create a dual-representation model, complemented by adaptive timestep scaling and temporal smoothing that enhanced stability.

GauMesh [23] proposed a framework that integrated triangle meshes within a single differentiable pipeline. It leveraged tracked meshes for representing smooth surfaces and detailed textures, while using 3D Gaussians, deformed via a grid-based feature volume, to capture complex geometry, performing alpha-blending between the two representations for the final rendering.

Yang et al. [26] pioneered a technique treating time as an integral dimension, designing specialized 4D parameters. Duan et al. [5] developed a rotor-based 4D representation for complex space-time transformations. Addressing the significant storage demands associated with direct 4D Gaussian optimization, 4D Scaffold GS [4] extends 3D scaffolding into 4D space. It utilizes sparse 4D grid-aligned anchors, each holding a compressed feature vector, to model sets of neural 4D Gaussians representing local spatio-temporal regions.

## 3. Preliminary: 3D Gaussian Splatting

3DGS [11] represents scenes using explicit 3D Gaussian primitives, enabling efficient novel view synthesis. Each Gaussian is parameterized by: (1) center position  $\mu \in \mathbb{R}^3$ , (2) covariance matrix  $\Sigma \in \mathbb{R}^{3 \times 3}$  decomposed as  $\Sigma = R S S^T R^T$  (where  $R$  is rotation from quaternion  $\mathbf{q}$  and  $S$  is diagonal scaling), (3) view-dependent color via Spherical Harmonics, and (4) opacity  $\alpha \in [0, 1]$ .

For rendering, 3D Gaussians are projected to 2D via transformation matrix  $W$  with Jacobian  $J$ , yielding 2D covariance  $\Sigma' = J W \Sigma W^T J^T$ . The final pixel color is computed through alpha-blending:  $C(p) = \sum_{i=1}^N \alpha_i c_i \prod_{j=1}^{i-1} (1 - \alpha_j)$ . Parameters are optimized by minimizing  $\mathcal{L} = (1 - \lambda) \mathcal{L}_1 + \lambda \mathcal{L}_{\text{D-SSIM}}$  between rendered

and ground-truth images. Adaptive density control refines the representation through pruning low-opacity Gaussians and densification via cloning or splitting based on gradients.

## 4. Methodological Analysis

Following the categorization in Section 1, we examine the key methodological adaptations applied to the foundational 3DGS framework for dynamic scene reconstruction, focusing on the distinguishing characteristics of each paradigm.

### 4.1. Structure-Guided

Structure-guided methods maintain canonical 3D Gaussians and predict temporal transformations via auxiliary structures. These methods introduce specialized loss functions and regularization to ensure geometric plausibility: SC-GS [9] employs As-Rigid-As-Possible (ARAP) loss to encourage locally rigid transformations, while DynMF [12] uses L1 sparsity regularization and rigid loss for coherent motion trajectories. Training typically involves jointly optimizing Gaussian parameters and deformation network weights, often with annealing schedules to stabilize convergence [24]. Timestamps are embedded (via sinusoidal or learned encodings) and fed to deformation networks alongside spatial inputs. Some methods bypass neural networks, such as Gaussian-Flow [15], by optimizing time-domain polynomials and Fourier components directly for faster convergence with fewer parameters.

### 4.2. Gaussian-Centric

Gaussian-centric methods encode temporal dynamics directly into primitives, extending 3D Gaussians to 4D or parameterizing attributes with continuous functions. This introduces new regularization requirements: 4D-Rotor GS [5] enforces 4D spatiotemporal consistency loss and entropy loss for sharper opacity, while STG [13] uses temporal radial basis functions requiring additional smoothness constraints. Optimization involves 4D transformation matrices with time-dependent parameters, often compressed via functional representations [10, 15]. Training strategies include temporal slicing—batching samples from adjacent time intervals to leverage local continuity [5]—and temporal weighting mechanisms that focus gradients on relevant contributions [26]. Advanced systems employ time-conditioned multi-pass optimization, accumulating losses across temporal slices before gradient updates to ensure consistent 4D trajectories [5].

## 5. Experiments

We evaluate representative methods from both paradigms on the D-NeRF dataset [19], a widely-used benchmark comprising eight synthetic scenes with animated objects exhibiting large motions and non-rigid deformations. We se-

lected methods with publicly available codebases from each category and established static 3DGS [11] as a baseline. All experiments were conducted on a NVIDIA RTX 4090 GPU.

We report Peak Signal-to-Noise Ratio (PSNR) [8], Structural Similarity Index Measure (SSIM) [21], and Learned Perceptual Image Patch Similarity (LPIPS) [28] for reconstruction quality, alongside frames per second (FPS), training time, and model size for computational performance.

### 5.1. Reconstruction Quality

All dynamic methods substantially outperform the static 3DGS baseline, confirming the necessity of temporal modeling. Structure-guided methods demonstrate superior reconstruction fidelity, with SC-GS [9] leading across all metrics, followed by Deformable 3DGS [24] and GaGS [16]. Gaussian-centric methods show wider variation, with top performers [5, 13, 26] achieving competitive but generally lower quality than leading structure-guided approaches.

### 5.2. Computational Cost

Gaussian-centric methods achieve significantly higher rendering speeds, with 4D Rotor GS [5] reaching 1257 FPS and several others [4, 13, 26] exceeding structure-guided methods. Structure-guided rendering speeds vary considerably (12-192 FPS). The most significant difference lies in model size: structure-guided methods maintain compact representations (14-48 MB), while gaussian-centric methods vary widely, with some compact [4] but others exceeding 1 GB [17, 26].

## 6. Discussion

A key finding is the consistent trend towards higher reconstruction fidelity among structure-guided methods, particularly evident in the monocular D-NeRF setup. Methods like SC-GS [9], which achieved a leading PSNR of 43.31 dB, along with strong performers like Deformable 3DGS [24] and GaGS [16], effectively leverage auxiliary structures. These include MLP-based deformation fields acting on canonical spaces [24], sparse control points driving deformation [26], or sparse convolutions for geometry-aware feature extraction [16], to impose strong spatio-temporal priors. Incorporating such structured information and associated priors, such as ARAP losses or multi-scale geometric features, appears crucial for enhancing temporal consistency and reconstruction reliability, especially given the under-constrained nature of learning from limited dynamic observations. Furthermore, by separating static representation (canonical Gaussians) from dynamic transformation (learned fields or control points), this paradigm generally maintains relatively compact model storage requirements.

Conversely, the gaussian-centric paradigm frequently excels in computational efficiency, particularly rendering speed. Methods encoding temporal changes directly into

Table 1. Reconstruction Quality and Computational Performance on D-NeRF [19] dataset. **Bold** denotes best performance and underline denotes second place.

Methods	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	FPS $\uparrow$	Train (min) $\downarrow$	Size (MB) $\downarrow$
<i>Baseline</i>						
3DGS [11]	23.19	0.930	0.080	170	10	<b>10</b>
<i>Structure-Guided</i>						
SC-GS [9]	<b>43.31</b>	<b>0.997</b>	<b>0.007</b>	123	17	22
Deformable 3DGS [25]	<u>39.51</u>	<u>0.991</u>	0.012	85	26	<u>14</u>
GaGS [16]	37.36	0.988	<u>0.010</u>	12	15	48
DynMF [12]	36.89	0.983	0.019	192	20	19
GauFre [14]	34.80	0.982	0.021	112	13	21
4DGS [22]	34.05	0.980	0.025	104	13	17
<i>Gaussian-Centric</i>						
STG [13]	34.27	0.986	0.032	276	23	41
4D Rotor GS [5]	34.26	0.970	0.033	<b>1257</b>	<b>8</b>	36
RealTime4DGS [27]	34.09	0.979	0.021	231	10	1432
4D Scaffold GS [4]	33.57	0.966	0.025	<u>293</u>	34	24
Compact 3DGS [10]	32.19	0.970	0.040	210	<u>9</u>	159
Dynamic 3DGS [17]	27.69	0.951	0.057	74	16	3014

Gaussian parameters, often using continuous functions like Fourier series [10], polynomials [15], or Radial Basis Functions [13], can bypass the need to query potentially large deformation networks at render time. The extension to true 4D primitives, as seen in RealTime4DGS [26], 4DGS [22], and notably 4D-Rotor GS [5], pushes rendering speeds to exceptional levels (Table 1), significantly exceeding both static 3DGS and structure-guided methods. However, this paradigm shows considerable variability in storage demands. While techniques like 4D Scaffold GS [4] achieve compactness through anchored representations, others that explicitly store 4D parameters or require per-frame optimization like Dynamic 3DGS [17] can lead to substantial model sizes, posing scalability challenges for long-duration videos or high resolution videos.

This comparison highlights a crucial trade-off in dynamic 3DGS: structure-guided methods often prioritize reconstruction fidelity and model compactness, potentially sacrificing some rendering speed, whereas gaussian-centric methods prioritize real-time rendering, sometimes at the expense of quality or storage efficiency. The choice between paradigms depends heavily on the specific application requirements: interactive AR/VR might favor the speed of gaussian-centric methods, while high-fidelity content creation might lean towards the quality offered by structure-guided approaches.

The importance of regularization is evident across both paradigms. The under-constrained nature of dynamic reconstruction necessitates techniques to prevent overfitting, temporal jitter, and implausible deformations. Our analysis highlights the effectiveness of various strategies, including ARAP losses [26], sparsity constraints [12], local rigidity and smoothness priors [17], and spatio-temporal consistency or entropy losses for 4D representations [5]. These

regularizers are critical for achieving stable training and temporally coherent results.

Building on these findings, several future directions emerge. Addressing large-scale unbounded dynamic scenes requires advancements in model scalability and adaptability to handle more intense motions and complex environmental changes, pushing beyond current benchmark limitations. This necessitates improving scene decomposition strategies, developing more efficient dynamic representations, exploring faster training techniques, and incorporating better geometric priors. Furthermore, research into novel data structures and primitives is crucial for better representing complex geometries, handling acute temporal events like appearances/disappearances, and optimizing resource usage. Exploring hybrid paradigms, perhaps using structure-guidance to regularize efficient gaussian-centric representations, could offer a compelling balance. Developing adaptive methods that select or switch representations based on scene complexity, performing thorough real-world validation on diverse captured data, integrating multi-modal priors like depth or flow, and designing more lightweight regularization techniques are all vital steps forward.

## 7. Conclusion

Dynamic 3DGS presents a vibrant research area characterized by a fundamental trade-off between reconstruction fidelity and computational efficiency. The structure-guided and gaussian-centric paradigms offer distinct solutions, each with unique advantages and challenges. Continued innovation in representations, regularization, and hybrid approaches will be essential to overcome current limitations and unlock the full potential of this technology for creating truly interactive and realistic dynamic digital worlds.

## References

- [1] Jeongmin Bae, Seoha Kim, Youngsik Yun, Hahyun Lee, Gun Bang, and Youngjung Uh. Per-gaussian embedding-based deformation for deformable 3d gaussian splatting. In *European Conference on Computer Vision (ECCV)*, 2024. 2
- [2] Ang Cao and Justin Johnson. Hexplane: A fast representation for dynamic scenes. *CVPR*, 2023. 1, 2
- [3] Guikun Chen and Wenguan Wang. A Survey on 3D Gaussian Splatting, 2024. arXiv:2401.03890 [cs]. 1
- [4] Woong Oh Cho, In Cho, Seoha Kim, Jeongmin Bae, Youngjung Uh, and Seon Joo Kim. 4D Scaffold Gaussian Splatting for Memory Efficient Dynamic Scene Reconstruction, 2024. arXiv:2411.17044 [cs]. 2, 3, 4
- [5] Yuanxing Duan, Fangyin Wei, Qiyu Dai, Yuhang He, Wenzheng Chen, and Baoquan Chen. 4d-rotor gaussian splatting: Towards efficient novel-view synthesis for dynamic scenes. In *Proc. SIGGRAPH*, 2024. 1, 2, 3, 4
- [6] Bardienus P Duisterhof, Zhao Mandi, Yunchao Yao, Jia-Wei Liu, Jenny Seidenschwarz, Mike Zheng Shou, Ramanan Deva, Shuran Song, Stan Birchfield, Bowen Wen, and Jeffrey Ichnowski. Deformgs: Scene flow in highly deformable scenes for deformable object manipulation. *WAFR*, 2024. 2
- [7] Zhiyang Guo, Wengang Zhou, Li Li, Min Wang, and Houqiang Li. Motion-aware 3D gaussian splatting for efficient dynamic scene reconstruction. *arXiv preprint arXiv:2403.11447*, 2024. 2
- [8] Alain Horé and Djemel Ziou. Image Quality Metrics: PSNR vs. SSIM. In *2010 20th International Conference on Pattern Recognition*, pages 2366–2369, 2010. ISSN: 1051-4651. 3
- [9] Yi-Hua Huang, Yang-Tian Sun, Ziyi Yang, Xiaoyang Lyu, Yan-Pei Cao, and Xiaojuan Qi. Sc-gs: Sparse-controlled gaussian splatting for editable dynamic scenes. *arXiv preprint arXiv:2312.14937*, 2023. 2, 3, 4
- [10] Kai Katsumata, Duc Minh Vo, and Hideki Nakayama. A Compact Dynamic 3D Gaussian Representation for Real-Time Dynamic View Synthesis, 2024. arXiv:2311.12897 [cs]. 1, 2, 3, 4
- [11] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 1, 2, 3, 4
- [12] Agelos Kratimenos, Jiahui Lei, and Kostas Daniilidis. Dynmf: Neural motion factorization for real-time dynamic view synthesis with 3d gaussian splatting. *ECCV*, 2024. 2, 3, 4
- [13] Zhan Li, Zhang Chen, Zhong Li, and Yi Xu. Spacetime gaussian feature splatting for real-time dynamic view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8508–8520, 2024. 1, 2, 3, 4
- [14] Yiqing Liang, Numair Khan, Zhengqin Li, Thu Nguyen-Phuoc, Douglas Lanman, James Tompkin, and Lei Xiao. Gaufré: Gaussian deformation fields for real-time dynamic novel view synthesis. In *Proc. IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2025. 2, 4
- [15] Youtian Lin, Zuozhuo Dai, Siyu Zhu, and Yao Yao. Gaussian-flow: 4d reconstruction with dynamic 3d gaussian particle. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21136–21145, 2024. 1, 2, 3, 4
- [16] Zhicheng Lu, Xiang Guo, Le Hui, Tianrui Chen, Ming Yang, Xiao Tang, Feng Zhu, and Yuchao Dai. 3d geometry-aware deformable gaussian splatting for dynamic view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2, 3, 4
- [17] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *3DV*, 2024. 2, 3, 4
- [18] Keunhong Park, Utkarsh Sinha, Jonathan T. Barron, Sofien Bouaziz, Dan B Goldman, Steven M. Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. *ICCV*, 2021. 1
- [19] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. *arXiv preprint arXiv:2011.13961*, 2020. 1, 3, 4
- [20] Qianqian Wang, Vickie Ye, Hang Gao, Weijia Zeng, Jake Austin, Zhengqi Li, and Angjoo Kanazawa. Shape of motion: 4d reconstruction from a single video. In *International Conference on Computer Vision (ICCV)*, 2025. 2
- [21] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004. Conference Name: IEEE Transactions on Image Processing. 3
- [22] GuanJun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 20310–20320, 2024. 2, 4
- [23] Yuting Xiao, Xuan Wang, Jiafei Li, Hongrui Cai, Yanbo Fan, Nan Xue, Minghui Yang, Yujun Shen, and Shenghua Gao. Bridging 3d gaussian and mesh for freeview video rendering. *CoRR*, abs/2403.11453, 2024. 2
- [24] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. *arXiv preprint arXiv:2309.13101*, 2023. 1, 2, 3
- [25] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3D Gaussians for High-Fidelity Monocular Dynamic Scene Reconstruction, 2023. arXiv:2309.13101 [cs]. 4
- [26] Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. In *International Conference on Learning Representations (ICLR)*, 2024. 1, 2, 3, 4
- [27] Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time Photorealistic Dynamic Scene Representation and Rendering with 4D Gaussian Splatting, 2024. arXiv:2310.10642 [cs]. 4
- [28] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unreasonable Effectiveness of

Deep Features as a Perceptual Metric. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. ISSN: 2575-7075. [3](#)