Beyond Static Perception: Integrating Temporal Context into VLMs for Cloth Folding

Oriol Barbany¹ Adrià Colomé¹ Carme Torras¹ ¹Institut de Robòtica i Informàtica Industrial, CSIC-UPC {obarbany, acolome, torras}@iri.upc.edu

https://barbany.github.io/bifold



Fig. 1: **Motivation**: In multi-step tasks like cloth folding, the state can become visually ambiguous due to self-occlusions or compounding errors from past actions. Relying solely on the current observation can make the next action ill-defined (left). By incorporating temporal context (right), the model gains critical information enabling more accurate perception and decision-making.

Abstract—Manipulating clothes is challenging due to their complex dynamics, high deformability, and frequent selfocclusions. Garments exhibit a nearly infinite number of configurations, making explicit state representations difficult to define. In this paper, we analyze *BiFold* [1], a model that predicts language-conditioned pick-and-place actions from visual observations, while implicitly encoding garment state through end-to-end learning. To address scenarios such as crumpled garments or recovery from failed manipulations, BiFold leverages temporal context to improve state estimation. We examine the internal representations of the model and present evidence that its fine-tuning and temporal context enable effective alignment between text and image regions, as well as temporal consistency.

Index Terms—VLM, Robotic Cloth Folding, Deformable Object Manipulation, LoRA Fine-Tuning, Temporal Consistency

I. INTRODUCTION

Cloth folding is a fundamental and repetitive task in daily life, present in both domestic and industrial settings. It is a labor-intensive activity that consumes significant time and effort and would greatly benefit from automation. Robotic cloth folding has the potential to free humans from these tedious chores and increase the productivity in sectors such as laundry services, textile manufacturing, and even in-home robotics. However, achieving robust and versatile robotic cloth folding is challenging due to the inherent complexity of cloth manipulation, including its high deformability, self-occlusions, and diverse material properties, topologies, and textures.

While some robotized tasks can be performed in open-loop or with simple sensing systems, perception plays a key role in deformable object manipulation. As in other vision-based tasks, changes in color distribution due to varying lighting conditions or textures can cause system failures. A standard solution in prior work on cloth manipulation is to discard color information and rely solely on depth images [2], [3], [4], [5]. However, depth images obtained from real sensors are noisy and color provides useful state recognition cues.

Most previous works decouple perception from the manipulation policy by explicitly extracting a representation and predicting actions based on it. In most cases, this representation is mesh-based [6], [7], [8], [9], [10], but alternative approaches include explicit surfaces [11], Gaussians [12], [13], implicit surfaces [14], [15], and keypoints [16], among others. Nonetheless, each has limitations, whether in flexibility to new topologies, resolution for capturing wrinkles or crumpled states, or richness of encoded information for effective manipulation.

In this work, we build on the BiFold model [1], which adapts a pretrained Vision Language Model (VLM) for language-conditioned bimanual cloth folding. Rather than introducing a new method, this workshop paper presents a deeper analysis of the internal mechanisms learned by BiFold. We explore how the model encodes visual and linguistic information across time, focusing on feature representations, attention patterns, and the role of temporal context. Our goal is to understand how BiFold interprets complex cloth states and leverages history for accurate perception, offering insights that can inform future work in deformable object manipulation.

II. METHODOLOGY

We focus on bimanual cloth folding starting from flattened configurations with sub-goals specified through natural language instructions. Specifically, we analyze the BiFold model [1], which we introduce below. For further details on the model architecture, dataset creation, and comprehensive experimental results, we refer the reader to the original paper.



Fig. 2: **BiFold architecture:** BiFold fine-tunes a pre-trained SigLIP [17] with LoRA [18] to obtain image and text features. The same image encoder is used for the current and past observations. Each token sequence is prepended with a different learned token, and we add information about the modality and the position inside the sequence by adding a learned embedding and sinusoidal positional encodings. All tokens are concatenated and processed using a transformer encoder [19] and the output tokens of the current observation are decoded into probability distributions on the pixel space using convolutional networks.

At each time step t, given an RGB image observation \mathbf{o}_t and a language instruction ℓ_t , BiFold predicts an action \mathbf{a}_t using a learned policy $\pi_{\theta}(\mathbf{a}_t | \ell_t, \mathbf{o}_t, \mathbf{o}_{t-1}, \mathbf{o}_{t-2}, \dots, \mathbf{o}_{t-H})$ parametrized by θ . In the base configuration, BiFold sets H = 3, which we use in our experiments. Crucially, the past observations correspond to *keyframes*, not raw camera frames as in works like RT-1 [20]. Since consecutive frames captured at high frequency are often highly redundant, BiFold instead uses observations sampled before the execution of each action, which provide more meaningful temporal context.

To effectively learn from vision and language inputs in a low-data regime, BiFold adapts SigLIP [17], a contrastive VLM similar to CLIP [21] composed of independent transformer-based encoders for images and text that produce aligned features. SigLIP offers high-level semantic representations that are well-suited for language-conditioned action prediction [22], and has demonstrated a low simto-real gap in robotic applications [23]. SigLIP is adapted using parameter-efficient fine-tuning to learn task-specific features while preserving the generalization capabilities of the foundational model. Specifically, BiFold employs Low-Rank Adaptation (LoRA) [18], a widely adopted method for adapting transformer-based models to downstream tasks [24], [25], [26], [22]. This design choice reduces the number of trainable parameters to just 0.29% of the more than 200 million parameters in the vision and text encoders.

Given the token sequences for the current image, language instruction, and context images, BiFold prepends each sequence with a distinct learned embedding. To encode positional information, BiFold applies 1D and 2D sinusoidal encodings for text and image tokens, respectively. Additionally, it includes a learned modality embedding to further distinguish between text and image inputs. BiFold then fuses all tokens by concatenating the sequences and applying cross-modal attention layers, following a strategy similar to [27], [28].

The output tokens corresponding to the current observation are then fed into convolutional decoders, which produce probability distributions over the pixel space. BiFold finally samples the pick and place positions for both the left and right arms from the predicted distribution to form the action \mathbf{a}_t , which can be passed to robot-specific controllers to execute actions.

The policy parameters θ are learned using the BiFold dataset [1], which contains language-aligned bimanual folding actions across various garment types. The dataset is constructed through a novel pipeline that extracts discrete actions from the continuous human demonstrations in the VR-Folding dataset [6]. Each garment is then canonically aligned within its category, enabling the generation of language instructions based on positions in this normalized space. Notably, the entire procedure is fully automatic, allowing the dataset to scale efficiently to new demonstrations, garment assets, and even unseen categories. More details in the full paper [1].

BiFold minimizes the binary cross-entropy loss between the predicted probability distribution and a Gaussian Mixture Model (GMM), where each Gaussian is centered at one of the ground-truth pick or place positions and assigned equal weight. For the pick positions, we restrict the Gaussian's support to the segmentation mask to avoid grasps at positions where there's no cloth. The resulting target mixture is normalized to ensure it defines a valid probability distribution over the pixel space.

III. RESULTS

We conduct experiments and provide visualizations that offer insight into the internal behavior of BiFold to support the following claims:

- (i) Parameter-efficient fine-tuning of VLMs enables extracting more task-specific representations while retaining their robustness and open-vocabulary understanding.
- (ii) Incorporating a *keyframe* history improves the model's ability to recognize and interpret complex object states.

A. Quantitative comparison

We evaluate the performance of BiFold and two variants of it: without context and with context using the consecutive frames rather than the previous *keyframes* to evaluate (ii).

In Table I, we report the Average Precision (AP) at different pixel thresholds and the output quantile, which indicates how likely are the ground truth positions in the predicted pick and

TABLE I: Quantitative metrics on BiFold dataset [1]



Observations $(\mathbf{o}_{t-2}, \mathbf{o}_{t-1}, \text{ and } \mathbf{o}_t)$

Fig. 3: Image encoder features: While pre-trained SigLIP features offer aligned language and visual embeddings, they struggle to capture semantic correspondences across time steps. In contrast, BiFold's fine-tuned version focuses more effectively on the manipulated cloth, enhancing the model's ability to distinguish task-relevant regions critical for folding.

place distributions. We can see that BiFold performs best, and its version with consecutive frames achieves, in most cases, worse metrics than using no context at all. We hypothesize this is due to considering uninformative tokens, which can make the model more susceptible to noise and harder to train effectively. Including all past frames introduces redundant or irrelevant visual information, especially in high-frame-rate settings, leading to feature dilution and potentially misaligned attention. Moreover, the increased sequence length demands more computational resources and may hinder learning, as the model needs to distinguish useful signals from a larger pool of mixed-quality tokens.

B. Implicit feature matching

To qualitatively assess (i), we analyze the patch-level features extracted by the image and text encoders. Specifically, we apply Principal Component Analysis (PCA) and visualize the first three components as RGB color maps. Following DINOv2 [29], we retain only those patches whose first principal component is non-negative To address (ii) and evaluate the impact of temporal context, we perform PCA on the aggregated features from both the current observation and its history of keyframes.



Fig. 4: Text-Image attention: Attention scores from text to image patches to the current observation and keyframes.

As shown in Fig. 3, the pre-trained features retained after PCA thresholding are not well localized on the garment, and there is a lack of alignment across frames. In contrast, after fine-tuning, the features more effectively isolate the main object from the background, and exhibit semantically meaningful correspondences across frames-such as consistent activation on sleeves, the center of the shirt, and their contours.

C. Attention maps

Before the transformer encoder in Fig. 2, the features for the input text and each of the observations are extracted independently. We now analyze how text tokens attend to visual tokens within the transformer encoder, providing insight into how language guides perception. In Fig. 4, we visualize attention maps over the current observation and all previous keyframes (when available) for task-relevant text tokens particularly informative for cloth folding. The token "bottom" attends to the lower regions of garments, while "fold" exhibits bimodal attention corresponding to plausible bimanual pick locations. In both cases, the attention maps reveal consistent spatial focus across temporal frames, supporting our claim (ii) that historical context enhances state recognition.

IV. DISCUSSION

Given the high degrees of freedom of garments and the challenges in explicitly modeling them, BiFold [1] implicitly encodes the state using a fine-tuned VLM. We showed that BiFold learns attention patterns aligned with the folding task and focuses on relevant features, effectively segmenting the manipulated cloth without supervision as in [29]. Although Bi-Fold learns cross-frame mappings without explicit supervision, additional learning signals such as NOCS [30] images could further encourage temporal consistency and spatial grounding.

While selecting keyframes is straightforward in our setting, this becomes more challenging in manipulation tasks where actions are not easily segmentable or involve many interdependent steps, making intermediate frames potentially valuable. One solution could be to automatically select the most informative frames from a history buffer or allocate tokens proportionally to their importance [31].

REFERENCES

- O. Barbany, A. Colomé, and C. Torras, "BiFold: Bimanual Cloth Folding with Language Guidance," arXiv:2501.16458, 2025.
- [2] W. Wang, G. Li, M. Zamora, and S. Coros, "TRTM: Template-based Reconstruction and Target-oriented Manipulation of Crumpled Cloths," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [3] Y. Deng, K. Mo, C. Xia, and X. Wang, "Learning Language-Conditioned Deformable Object Manipulation with Graph Dynamics," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [4] K. Mo, C. Xia, X. Wang, Y. Deng, X. Gao, and B. Liang, "Foldsformer: Learning Sequential Multi-Step Cloth Manipulation With Space-Time Attention," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 760–767, 2023.
- [5] T. Weng, S. Bajracharya, Y. Wang, K. Agrawal, and D. Held, "Fabricflownet: Bimanual cloth manipulation with a flow-based policy," in *Conference on Robot Learning (CoRL)*, 2021.
- [6] H. Xue, W. Xu, J. Zhang, T. Tang, Y. Li, W. Du, R. Ye, and C. Lu, "GarmentTracking: Category-Level Garment Pose Tracking," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2023, pp. 21233–21242.
- [7] C. Chi and S. Song, "Garmentnets: Category-level pose estimation for garments via canonical space shape completion," in *IEEE International Conference on Computer Vision (ICCV)*, 2021.
- [8] E. Caldarelli, A. Colomé, C. Ocampo-Martinez, and C. Torras, "Quadratic dynamic matrix control for fast cloth manipulation," in 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2023, pp. 8178–8185.
- [9] Z. Huang, X. Lin, and D. Held, "Mesh-based dynamics with occlusion reasoning for cloth manipulation," in *Robotics: Science and Systems* (*RSS*), 2022.
- [10] X. Lin, Y. Wang, Z. Huang, and D. Held, "Learning visible connectivity dynamics for cloth smoothing," in *Conference on Robot Learning* (*CoRL*), 2022, pp. 256–266.
- [11] O. Barbany, A. Colomé, and C. Torras, "Deformable surface reconstruction via riemannian metric preservation," *Computer Vision and Image Understanding*, vol. 249, p. 104155, 2024.
- [12] A. Longhini, M. Büsching, B. P. Duisterhof, J. Lundell, J. Ichnowski, M. Björkman, and D. Kragic, "Cloth-Splatting: 3D State Estimation from RGB Supervision for Deformable Objects," in *Conference on Robot Learning (CoRL)*, 2024.
- [13] M. Zhang, K. Zhang, and Y. Li, "Dynamic 3d gaussian tracking for graph-based neural dynamics modeling," in *Conference on Robot Learning (CoRL)*, 2024.
- [14] R. Li, C. Dumery, Z. Deng, and P. Fua, "Reconstruction of Manipulated Garment with Guided Deformation Prior," in Advances in Neural Information Processing Systems (NeurIPS), 2024.
- [15] B. Guillard, F. Stella, and P. Fua, "Meshudf: Fast and differentiable meshing of unsigned distance field networks," in *European Conference* on Computer Vision (ECCV), 2022.
- [16] T. Lips, V.-L. De Gusseme, and F. Wyffels, "Learning keypoints for robotic cloth manipulation using synthetic data," *IEEE Robotics and Automation Letters*, vol. 9, no. 7, pp. 6528–6535, 2024.
- [17] X. Zhai, B. Mustafa, A. Kolesnikov, and L. Beyer, "Sigmoid loss for language image pre-training," in *IEEE International Conference on Computer Vision (ICCV)*, 2023, pp. 11975–11986.
- [18] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in *International Conference on Learning Representations (ICLR)*, 2022.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. u. Kaiser, and I. Polosukhin, "Attention is All you Need," in *Neural Information Processing Systems (NeurIPS)*, vol. 30, 2017.

- [20] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu, J. Ibarz, B. Ichter, A. Irpan, T. Jackson, S. Jesmonth, N. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, K.-H. Lee, S. Levine, Y. Lu, U. Malla, D. Manjunath, I. Mordatch, O. Nachum, C. Parada, J. Peralta, E. Perez, K. Pertsch, J. Quiambao, K. Rao, M. Ryoo, G. Salazar, P. Sanketi, K. Sayed, J. Singh, S. Sontakke, A. Stone, C. Tan, H. Tran, V. Vanhoucke, S. Vega, Q. Vuong, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich, "RT-1: Robotics Transformer for Real-World Control at Scale," arXiv:2212.06817, 2022.
- [21] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, "Learning Transferable Visual Models From Natural Language Supervision," in *International Conference on Machine Learning (ICML)*, 2021.
- [22] M. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn, "OpenVLA: An Open-Source Vision-Language-Action Model," *arXiv* preprint arXiv:2406.09246, 2024.
- [23] K. Ehsani, T. Gupta, R. Hendrix, J. Salvador, L. Weihs, K.-H. Zeng, K. P. Singh, Y. Kim, W. Han, A. Herrasti, R. Krishna, D. Schwenk, E. VanderBilt, and A. Kembhavi, "Spoc: Imitating shortest paths in simulation enables effective navigation and manipulation in the real world," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024, pp. 16238–16250.
- [24] M. Zanella and I. Ben Ayed, "Low-rank few-shot adaptation of visionlanguage models," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024, pp. 1593–1603.
- [25] O. Barbany, M. Huang, X. Zhu, and A. Dhua, "Leveraging Large Language Models for Multimodal Search," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2024.
- [26] Y. Feng, J. Han, Z. Yang, X. Yue, S. Levine, and J. Luo, "Reflective planning: Vision-language models for multi-stage long-horizon robotic manipulation," arXiv:2502.16707, 2025.
- [27] Octo Model Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, C. Xu, J. Luo, T. Kreiman, Y. Tan, L. Y. Chen, P. Sanketi, Q. Vuong, T. Xiao, D. Sadigh, C. Finn, and S. Levine, "Octo: An open-source generalist robot policy," in *Robotics: Science and Systems (RSS)*, Delft, Netherlands, 2024.
- [28] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choromanski, T. Ding, D. Driess, A. Dubey, C. Finn, P. Florence, C. Fu, M. G. Arenas, K. Gopalakrishnan, K. Han, K. Hausman, A. Herzog, J. Hsu, B. Ichter, A. Irpan, N. Joshi, R. Julian, D. Kalashnikov, Y. Kuang, I. Leal, L. Lee, T.-W. E. Lee, S. Levine, Y. Lu, H. Michalewski, I. Mordatch, K. Pertsch, K. Rao, K. Reymann, M. Ryoo, G. Salazar, P. Sanketi, P. Sermanet, J. Singh, A. Singh, R. Soricut, H. Tran, V. Vanhoucke, Q. Vuong, A. Wahid, S. Welker, P. Wohlhart, J. Wu, F. Xia, T. Xiao, P. Xu, S. Xu, T. Yu, and B. Zitkovich, "RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control," arXiv:2307.15818, 2023.
- [29] M. Oquab, T. Darcet, T. Moutakanni, H. V. Vo, M. Szafraniec, V. Khalidov, P. Fernandez, D. HAZIZA, F. Massa, A. El-Nouby, M. Assran, N. Ballas, W. Galuba, R. Howes, P.-Y. Huang, S.-W. Li, I. Misra, M. Rabbat, V. Sharma, G. Synnaeve, H. Xu, H. Jegou, J. Mairal, P. Labatut, A. Joulin, and P. Bojanowski, "DINOv2: Learning robust visual features without supervision," *Transactions on Machine Learning Research*, 2024.
- [30] H. Wang, S. Sridhar, J. Huang, J. Valentin, S. Song, and L. J. Guibas, "Normalized Object Coordinate Space for Category-Level 6D Object Pose and Size Estimation," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019.
- [31] W. Yan, V. Mnih, A. Faust, M. Zaharia, P. Abbeel, and H. Liu, "Elastictok: Adaptive tokenization for image and video," in *International Conference on Learning Representations (ICLR)*, 2025.