Videos to Forces: Physics Based Climbing Control From Online Videos

Anonymous authors
Paper under double-blind review

Abstract

Climbing motions are especially challenging to reconstruct from videos because of their dense, complex contacts. Existing video-based reconstructions recover kinematics but not the required dynamics; physics-based motion tracking with reinforcement learning (RL) offers an alternative, still relies on these imperfect kinematic references and requires the policy to discover viable contact forces through exploration. We reinterpret a fictitious force that prior methods use for partial assistance, instead using it to fully support the motion as an intermediate dynamic reference. Although residual assistance must ultimately be removed, this force reveals how much support must ultimately come from physical contacts, and we use it to densely guide the policy to learn to clime using only the physical contact forces. Our method reconstructs challenging climbing behaviors from video-derived motion references. Across diverse wall geometries and climbing techniques, it reduces penetration, improves motion quality, and makes policy learning more reliable and efficient than prior kinematic and physics-based approaches. To our knowledge, this is the first method to reconstruct complex human climbing motions from video using only physical contact forces at test time.

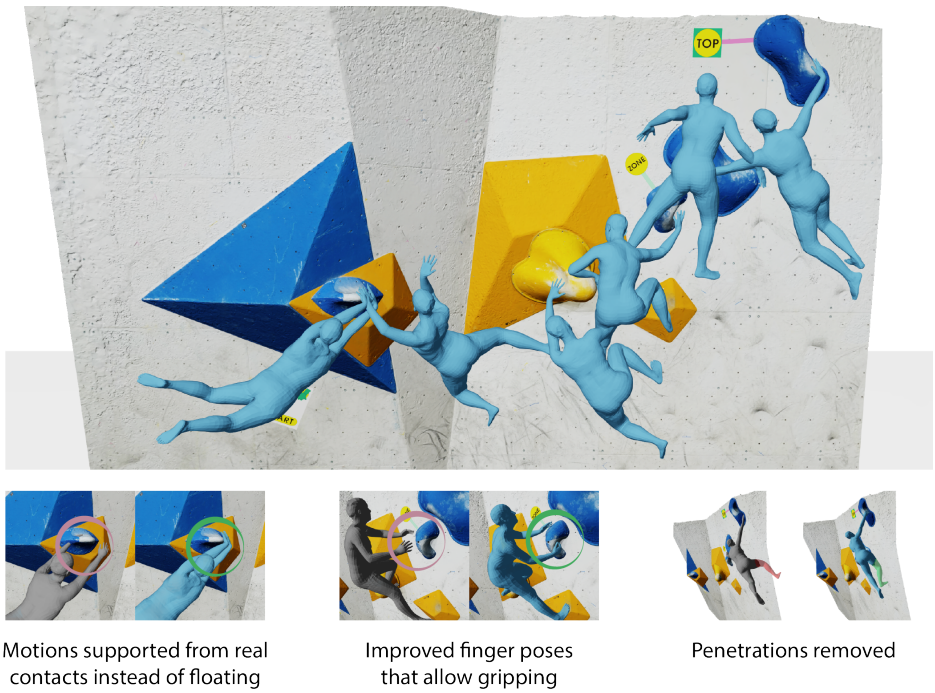

Figure 1: We present a method for recovering physically supported climbing motion from video. Here we show our method on a climb from the 2024 Boulder Japan Cup, Japan's national bouldering competition. Starting from the CoMotion (Newell et al., 2025) kinematic reconstruction (gray), our method produces natural, physically supported motion shown in blue. Here our method executes a dynamic start and finish together with a contact-rich middle sequence.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 Introduction

We consider the problem of reconstructing physically valid behavior from videos of humans performing complex contact interactions with the environment. In this work, we focus on climbing, a particularly challenging instance of this problem. Climbing requires dynamic whole-body motion, repeatedly making and breaking contact with irregular wall geometry while generating contact forces to support the body against gravity. Reconstructing such motions with physically consistent contact forces can help in realistic character animation for sports (Cha et al., 2015), as well as the analysis and teaching of athletic technique (Beltrán Beltrán et al., 2023; Blagus et al., 2024).

Recent foundation models for human pose estimation enable 4D reconstruction of kinematic human motion from video (Newell et al., 2025). Their outputs can be further refined with optimization-based methods that encourage geometric consistency and contact with the environment (Liu et al., 2026). However, these approaches still lack physical realism. As shown in Fig. 1, reconstructed motions can exhibit large penetrations with the wall geometry and unrealistic end-effector contacts. More fundamentally, they recover only kinematic trajectories, without the forces and dynamics required to explain how the motion is physically realized.

Physics-based motion tracking methods such as DeepMimic (Peng et al., 2018) train reinforcement learning (RL) policies to track kinematic references inside a physics simulator. They offer a way to obtain dynamically feasible motion, but they still rely on the imperfect kinematic references described above. The policy must therefore discover, through random exploration, how to correct these references and establish the physical contacts needed to support the motion. For contact-rich motions like climbing, this is particularly brittle, as small tracking errors can break contact and cause the character to fall under gravity. Prior work (Yuan & Kitani, 2020) alleviates this exploration problem using a "hand-of-God" residual force, but motions that depend on such forces are not fully supported by real contacts.

Our key insight is that although residual force must ultimately be removed, it is most useful when first applied in full. When fully supporting the motion, the residual force directly measures the support that physical contacts must eventually provide, turning an optimization aid into a dynamics signal that augments the usual kinematic reference. We use this signal in two complementary ways: first, to initialize a curriculum that gradually anneals residual forces and transfers them to physical contacts; and second, to derive per-contact force and torque targets that guide the policy towards the contact loading needed for such transfer.

To the best of our knowledge, we are the first to demonstrate learning of highly dynamic climbing motions in a physics-based simulator without residual forces at test time. Our curriculum progressively transfers support from fictitious residual forces to physical contacts, yielding motions that are fully contact-supported at test time. Compared with kinematic reconstruction baselines, our method reduces wall penetration and contact sliding while improving contact placement. Compared with standard physics-based tracking, our method better preserves pose quality as contact geometry becomes more challenging. Ablations further show that our dynamics guidance leads to more consistent training completion and stronger contact loading on difficult climbs.

2 Related Work

Contact-Aware Human Motion Reconstruction from Videos. Parametric body models such as SMPL (Loper et al., 2015) and MHR (Ferguson et al., 2025) represent human motion through articulated joint configurations and corresponding body mesh geometry. Large-scale datasets, ranging from clean motion-capture (Ionescu et al., 2013) to in-the-wild video with pseudo-labels (Kanazawa et al., 2019), have enabled training of foundation models for human motion reconstruction (Newell et al., 2025; Yang et al., 2026). To enforce contact-grounding, recent work jointly models humans with interacting objects (Cao et al., 2021) or surrounding scenes (Müller et al., 2025). Additionally, contact and non-penetration constraints can be applied (Hassan et al., 2019; Zhang et al., 2021). JOSH (Liu et al., 2026) extends this idea to joint optimization of human motion, scene geometry, and camera pose for human motions including climbing. However, these methods remain primarily *kinematic*: they recover the geometry of contact, but not the forces needed to physically realize the motion. We use such kinematic references as a starting point for learning physics-based controllers.

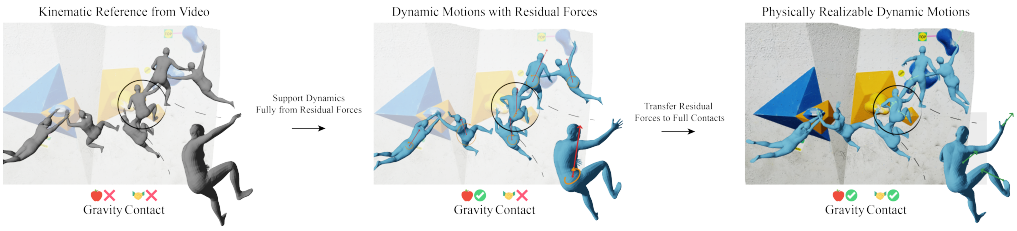

Figure 2: **Method Overview.** We first extract a kinematic reference from video, then learn a no-contact dynamic reference fully supported by residual forces, and finally train a contact-rich policy that transfers this support to physical contacts. Red, orange, and green arrows visualize residual force, residual torque, and contact forces, respectively. Closeups of the same representative frame reveal changes that are subtle from the distant view. The characters gain progressively better wall alignment from the raw reconstruction to the dynamic reference, followed by pose adjustments in the final policy that enable physical contact loading.

Physics-Based Human Motion Imitation. DeepMimic (Peng et al., 2018) established reinforcement learning (RL) as an effective method for converting kinematic motion references into physics-based controllers. When clean motion capture data is available, this approach can learn highly agile motion (Peng et al., 2021; Liao et al., 2025; Peng et al., 2022). Such approaches have also been applied to noisier data sources such as monocular videos (Shimada et al., 2020; Yuan et al., 2021), broadcast footage (Zhang et al., 2023) or imperfect references (Xu et al., 2025). However, climbing is substantially more challenging, requiring repeated contact switches on irregular surfaces while continuously supporting the body against gravity. In this work, we ease this difficult exploration problem by constructing an intermediate dynamic motion reference using residual forces.

Residual Forces for Physics-Based Control. Non-physical assistance has long been used to make difficult control problems easier to optimize, with external "hand-of-God" forces gradually removed as the controller improves (Van de Panne & Lamouret, 1995; Yu et al., 2018; Cao et al., 2025). More recent work applies similar ideas to contact-rich manipulation (Mandi et al., 2025; Zhao et al., 2025), using virtual object controllers during training. Residual Force Control (Yuan & Kitani, 2020) uses an fictitious root wrench to compensate for mismatch between reference and simulated dynamics, typically retaining it in the final controller. We instead use the residual force in two ways: as a curriculum, annealed gradually to zero, and as a dynamic reference, whose net external forces we redistribute into per-contact force and joint-torque targets.

3 Method

Our goal is to turn an imperfect kinematic climbing reference into motions supported entirely by contacts. Tracking such complex motions often requires residual forces to make exploration around imperfect references tractable. However, training such a policy still requires substantial exploration to discover viable contact strategies, because the reference specifies where the body should move, but not how contacts should support it. Our key observation is that if residual forces are first allowed to fully support the motions, they expose exactly this missing dynamics information. We therefore train our policy stage-wise: first extracting a kinematic reference, then learning a dynamic reference fully supported by residual forces, and finally using that reference to guide a contact-rich policy as residual support is transferred to physical contacts. Figure 2 summarizes this pipeline. Unless otherwise stated, all our policies are trained with PPO following the DeepMimic (Peng et al., 2018) tracking objective, with joint position targets as actions.

Extract kinematic references from videos. Our first goal is to extract kinematic reference motions that are approximately consistent with the ground-truth wall geometry. We apply an off-the-shelf model, CoMotion (Newell et al., 2025), to monocular climbing videos, obtaining a 4D human pose trajectory over time. Because this reconstruction has no access to wall geometry, we further refine the body scale, root pose, and joint rotations using a lightweight optimization. This optimization uses the video-derived motion as a prior while moving contacting end effectors towards their corresponding holds, reducing sliding during contact, aligning palm orientation with the contacted surface, and

preventing body-wall penetration. Full optimization details are provided in Section A.1 of the Appendix.

Extract dynamic references with Residual Force Control. Tracking the kinematic reference is difficult because climbing is supported by only a few contacts at a time, against complex wall geometries. For locomotion on the ground, a small error merely degrades the motion, which turns out to provide gradient signals for the policy to update its control. Climbing, however, involves an unforgiving reward landscape, where an incorrectly timed or placed contact can send the climber off the wall immediately, terminate the episode and stall reward signals. We therefore use Residual Force Control (Yuan & Kitani, 2020), augmenting the policy action with a fictitious 6D force and torque applied at the root, which we refer to as the residual force.

Rather than using this residual only to assist tracking, our goal in this stage is to use it to extract the dynamics required by the motion. We do so by disabling wall collision and allowing the residual force to fully support the character. The no-contact dynamics are therefore

$$M(q)\ddot{q} + h(q, \dot{q}) \;=\; S^\top \tau \;+\; B^\top F_{\text{res}}, \tag{1}$$

where q is the pose, M is the mass matrix, h collects Coriolis, centrifugal, and gravity terms, τ is the actuated joint torques produced by the policy, S selects all non-root joints, $F_{\text{res}} \in \mathbb{R}^6$ is the residual force applied at the root, and $B \in \mathbb{R}^{6\times n}$ selects the six floating-base degrees of freedom of the character.

In the final contact-rich policy, we want to replace F_{res} entirely with contact forces $F_c \in \mathbb{R}^3$ per end effector i. If the no-contact and contact-rich policies realize the same pose q, their dynamics must match:

$$F_{\text{res}} \;\equiv\; \textstyle\sum_{i=1}^{N_{ee}} (J_i^{\text{root}})^\top F_{c,i}, \quad \tau = \tau' - \sum_i (J_i^{\text{joint}})^\top F_{c,i} \tag{2}$$

where J_i^{root} and J_i^{joint} denote the root and actuated-joint components of the i-th end-effector Jacobian, τ is the joint torques applied in the final policy, τ' is the joint torques applied in the converged no-contact policy.

Thus, the residual force learned by the no-contact policy directly specifies the net wrench that physical contacts must ultimately reproduce, while the second relation specifies the corresponding change in joint torques. The full derivation is provided in Section A.2.

Generate physically realizable motion using dynamic references. Now that the no-contact policy has exposed the support required by the motion, our final stage is to transfer this support from residual forces to physical contacts. We cache the residual forces F_{res}^* at the convergence of the no-contact policy and use a curriculum to gradually convert F_{res}^* into contact forces F_c. The curriculum drives a residual scale $\alpha_k \in [0, 1]$ to zero over iteration k with step size Δ, $\alpha_0 = 1$, and maintains

$$\left\|F_{\text{res}}\right\| \le \alpha_k \left\|F^*\right\|, \quad \alpha_{k+1} \;=\; \begin{cases} \max(0,\ \alpha_k - \Delta)\,, & \text{if survival rate plateaus and pose reward above a threshold} \\ \alpha_k, & \text{otherwise} \end{cases} \tag{3}$$

Training is complete when $\alpha_k = 0$ so that the body weight is fully supported by wall contacts.

Additionally, unlike prior works that only use kinematic references as reward signals, we also use force references, specifically our cached residuals, to shape the policy's controls. We derive a loss guidance from equation 2, whose second half provides exactly the target joint torques of the contact-rich policy. These torques are valid targets for the policy as long as poses $q(t)$ are approximately maintained (see A.2 for more details).

To estimate joint torques, we first estimate contact forces from residual forces by solving a simple optimization, following the equivalence in the first half of equation 2. Denoting the force and torque components of the cached residual as $F_{\text{res}}^*(t) = [f^*(t), T^*(t)]$, we estimate $F_{c,i}(t)$ for each contacting end effector i in each frame t:

$$\begin{aligned} \min_{F_c} \quad & \sum_i \|F_{c,i}(t)\|^2 + \lambda \|F_{c,i}(t) - F_{c,i}(t-1)\|^2 \\ \text{s.t.} \quad & \sum_i F_{c,i}(t) = f^*(t), \quad \sum_i F_{c,i}(t) \times (p_i(t) - p_{\text{root}}(t)) = T^*(t), \\ & F_{c,i}(t) \cdot n_i(t) \ge f_{\min}, \quad \|F_{c,i}(t) - (F_{c,i}(t) \cdot n_i(t)) n_i(t)\| \le \mu (F_{c,i}(t) \cdot n_i(t)). \end{aligned} \tag{4}$$

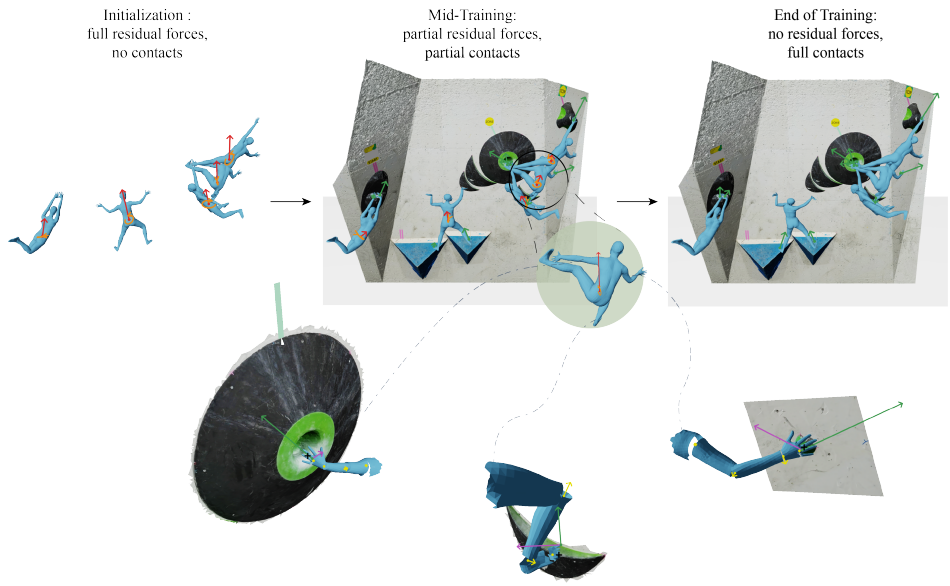

Figure 3: **Loss guidance from residual forces**. Top: progression from the residual-supported initialization to the final contact-supported policy. Red and orange arrows denote residual force and torque, while green arrows show physical contact forces. Bottom: close-ups illustrate how the residual-derived guidance guides the policy controls: yellow arrows indicate the induced joint corrections, and pink arrows show the resulting end-effector response. In the illustrated frame, the guidance encourages the left hand to press downward, the right hand to load the hold more strongly, and the right foot to step down, producing contact forces consistent with the required motion dynamics, all aligning with what the frame's motion requires.

where p_i and p_{root} are the positions of the end effectors and the root, n_i is the wall normal at the end effectors' contact point, and μ is the friction coefficient there. $f_{min} = 20N$ is a minimum force magnitude requirement for contacts.

Solving F_c from this optimization allows us to compute policy action targets a^{targ} using equation 2:

$$a^{\text{targ}} = [k_p^{-1}\tau^{\text{targ}}, \alpha F_{\text{res}}], \quad \text{where } \tau^{\text{targ}} = \tau' - \sum_i (J_i^{\text{joint}})^\top F_{c,i} \tag{5}$$

We compute the differences between these targets and the policy actions for a loss that we use alongside the PPO loss L_{ppo} to guide training:

$$L = L_{\text{ppo}} + w_{\text{targ}} L_{\text{targ}}, \quad L_{\text{targ}}(\theta) = \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \left\| \mu_\theta(s) - a^{\text{targ}}_{t(s)} \right\|^2 \tag{6}$$

where $\mathcal{S}$ is the set of rollout states. Figure 3 shows how this guidance term encourages joint rotations to drive the end effectors to press harder into directions that better support the motions, thereby creating more contact force loading to offset residuals.

4 Results

Figure 4 qualitatively demonstrates our method across diverse wall geometries and climbing styles. We will further evaluate on each part of our method: (1) we compare our full method against the same pipeline without residual-force guidance, measuring whether the dynamic guidance improves training completion and contact loading, particularly on challenging contact configurations; (2) we compare against DeepMimic (Peng et al., 2016) to examine how standard physics-based tracking behaves as contact geometry becomes less forgiving, including failures through pose drift, penetration, missed contacts and falling; (3) we examine how the curriculum transfers support from assistive forces to



Figure 4: **Our method across diverse climbing styles and contact regimes.** Each wall shows one motion segment, with lighter blue poses tracing the past trajectory leading to the current darker blue pose. From left to right: a power move supported primarily by the two arms; a controlled hand reach maintaining three to four contacts throughout; a highly dynamic jumping sequence with rapid contact changes; and a slow balance move supported mainly through both feet.

physical contacts; (4) we evaluate how the quality of the kinematic reference affects downstream learning; (5) we additionally provide a use of our method to non-climbing motions. In a simpler cartwheel motion, we show that both our method and DeepMimic produce matching quality.

Setup. All our policies are trained in Isaac Sim with an SMPL character. Each variant for the same motion is trained on the same reference, with the same character file and collisions enabled.

Effect of residual-force guidance on policy learning Figure 5 shows DeepMimic vs. our method without guidance vs. our method with guidance. Each is evaluated on the following metrics: (1) success: the share of 512 rollouts that reach the end of the motion without falling, (2) contact loading: percentage of frames where the character correctly loads all end effectors that are supposed to be in contact, (3) penetration: maximum depth of the end effectors into the wall, (4) joint angle error: the root mean square (RMS) against the reference. Our guided policy finishes training within the allotted time on every motion, while both the unguided variant and DeepMimic time out on some climbs, leading to the low success rates. On easier motions with smooth contact geometry and a relatively constrained way to distribute support, guidance achieves contact loading comparable to the unguided variant. On harder motions with more irregular geometry or multiple plausible contact-force distributions, guidance produces stronger and more consistent contact loading. The comparison between our guided and unguided variants isolates the effect of the dynamics guidance. While DeepMimic has some reasonable success rates, it trades penetration and pose fidelity for survival; we examine this behavior qualitatively in Figure 6.

Comparison with standard physics-based motion tracking under challenging contact geometry. Figure 6 examines how DeepMimic behaves as climbing wall surface becomes less smooth. On large, smooth holds, it can it can track the reference reasonably well. However, with smaller holds or holds with more complex geometry, the policy may fail to establish the intended contact or compensate through pose distortions and penetration. These failure modes are consistent with the fact that the policy receives no dynamic support or guidance on how each contact should support the motion, and can settle into local minima of successful motions that are unnatural or physically unrealizable. In contrast, our method completes all motions with natural poses.

Transfer of support through the residual-force curriculum. Figure 7 shows how our curriculum progressively transfers residual forces to contacts. Early in training, the motion relies almost entirely on the residual force. As training progresses, residual support decreases to zero while contact support increases to approximately the cached residual levels.

Comparison with kinematic reference reconstruction. Since our policy depends on the quality of its kinematic reference, we compare our reference extraction against JOSH (Liu et al., 2026),

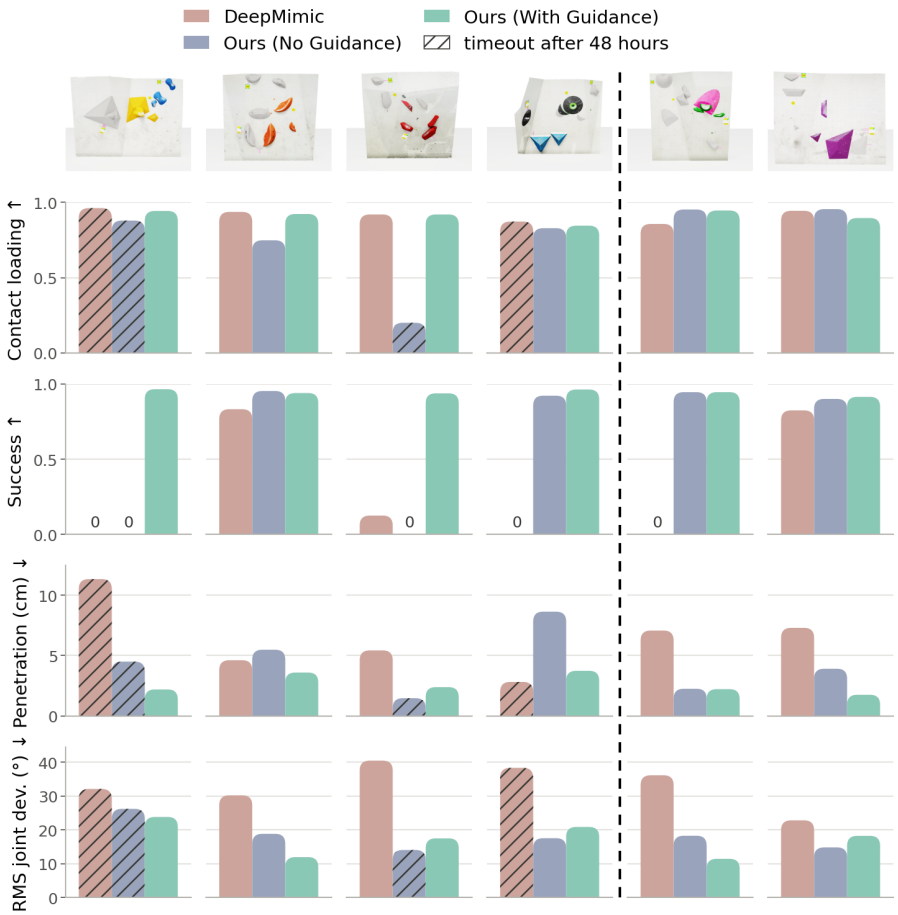

Figure 5: **Per-motion comparison of DeepMimic, our method without residual-force guidance, and our full method, all evaluated with zero residual force.** Motions are grouped into harder climbs with irregular contact geometry or multiple plausible contact-force distributions (left) and easier climbs with smoother geometry and more constrained contact loading (right).

Table 1: Comparison with JOSH on seven climbs. Each cell is JOSH / Ours; lower is better for every metric, better value in bold.

Metric \Climb	1	2	3	4	5	6	7
p95 penetration depth (cm)	40.8 / **9.3**	41.2 / **9.5**	69.0 / **9.5**	64.9 / **8.9**	120.0 / **9.4**	26.1 / **9.1**	98.8 / **9.3**
Contact-point speed during contact, mean (cm/s)	208.1 / **6.6**	137.1 / **4.7**	57.1 / **7.0**	65.3 / **6.0**	131.3 / **7.5**	55.4 / **10.0**	248.9 / **9.7**
Contact point to target, mean (cm)	90.8 / **8.4**	84.8 / **3.0**	67.3 / **7.9**	105.9 / **8.9**	127.8 / **5.8**	87.6 / **6.7**	175.7 / **6.4**
Contact point to nearest surface, mean (cm)	14.7 / **2.5**	10.2 / **1.0**	8.1 / **4.2**	16.4 / **3.1**	35.6 / **2.5**	8.0 / **2.7**	22.3 / **2.0**

a monocular human-scene reconstruction method demonstrated on climbing. Unlike our method, JOSH jointly reconstructs both human motion and an unknown scene, whereas we assume access to the ground-truth wall mesh. We therefore evaluate each method relative to its own scene, measuring end-effector penetration, distance to the intended contact, and contact-point speed during contact. Table 1 shows that our references achieve substantially smaller penetration, closer contact placement, and less contact sliding. We next examine how this reference quality affects downstream policy training.

Effect of kinematic reference quality on policy learning. while we allow inevitable errors in kinematic reference extraction, we still require the reference to stay in a reasonable range around the ground truth. To isolate the effect of our kinematic refinement, we compare against a root-only variant of the method that only optimizes root positions to be of the CoMotion reconstruction to be



Figure 6: **Qualitative comparison with DeepMimic as climbing contact geometry becomes more challenging from left to right.** On all three motions our method completes without pose drifting. **Left**: DeepMimic tracks the motion well on large, smooth holds, but fails once it reaches the later small-handhold contact; the red character marks the frame that rollout fails to reach, while the twisted left-arm pose shows pose drifting before failure. **Middle**: DeepMimic fails near the start when the right arm is blocked by the protruding pink hold; restarting from a later frame allows completion but introduces pose errors, including back-of-hand contact and a twisted right shoulder. **Right**: DeepMimic begins with the right arm deeply penetrating the wall and later fails again, shown by the red backward-flipping pose; restarting from a later frame recovers the remainder of the motion but still produces incorrect palm orientation.

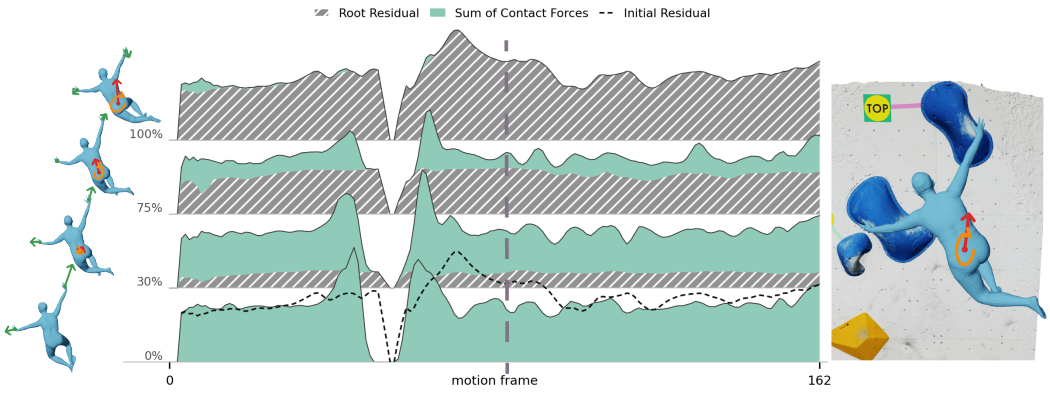

Figure 7: **Transfer of support from residual forces to physical contacts over training.** We show residual-force magnitude and the summed contact force projected onto the cached residual force direction throughout the curriculum. The dashed vertical line marks an example frame. Its state at the end of no-contact policy training is shown on the right; the sequence on the left shows the same frame over the contact-rich policy training.

close to the wall, without refining joint rotations or contacts. Table 2 shows that our refined references consistently reduce penetration and joint deviation, while showing more consistent success rate and contact loading.

Table 2: **Effect of kinematic reference quality on downstream policy learning.** Each cell reports *Root-only / Ours*, with both policies evaluated at zero residual force. Root-only applies only the root-position correction to CoMotion, while Ours additionally refines joint rotations and contacts. Higher is better for success and contact loading; lower is better for penetration and RMS joint deviation.

Climb	Success ↑	Contact loading ↑	Penetration, max (cm) ↓	RMS joint dev. (°) ↓
1	**0.99** / 0.95	**0.94** / 0.91	1.6 / **1.4**	28.9 / **16.5**
2	0.31 / **0.96**	0.40 / **0.97**	13.3 / **2.3**	33.8 / **19.1**
3	**0.96** / 0.93	0.88 / **0.90**	6.2 / **2.6**	25.8 / **21.7**
4	0.00 / **0.90**	0.60 / **0.96**	9.7 / **3.9**	15.5 / **14.8**

Application beyond climbing. Our framework is not specific to climbing and can be easily applied to any other motions. Many such motions involve simpler contact geometry and are therefore easier for standard physics-based tracking. We test this on a cartwheel motion, where our method and DeepMimic achieve comparable motion and contact quality. Together with Figure 6, this suggests that our method retains performance on simple contact behaviors while providing increasing benefits as contact geometry becomes more constrained and irregular.

5 Conclusion

We presented a physics-based framework for reconstructing climbing motions from video by using residual force as a dynamics signal in addition to the usual kinematic reference. By first allowing residual forces to fully support the motion, we obtain a measure of the support that contacts must eventually provide. We use this force to guide a curriculum that drives towards contact-supported motions. Across diverse climbing motions and wall geometries, our method reduces penetration, improves motion quality, and makes policy learning more reliable, while requiring no residual forces at test time. More broadly, we hope this work can serve as a step towards extracting richer dynamics signals from videos and using them to guide physics-based policy learning.

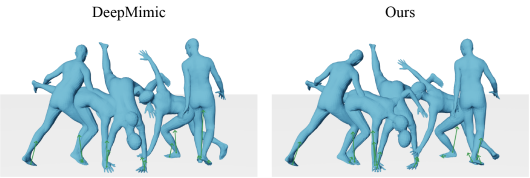

Figure 8: **Qualitative comparison on a cartwheel motion.** DeepMimic and our method both reproduce the motion with comparable pose and contact quality. Unlike climbing, this motion involves simple contacts on a continuous, smooth surface, where standard physics-based tracking is already effective.

6 Limitations

Dependence on accurate wall geometry. Our kinematic reference reference extraction assumes access to the ground truth wall mesh, which we use to anchor characters' root position, avoid penetration, and orient the end effectors relative to the wall geometry. In settings where the wall geometry must itself be reconstructed from video, errors in the scene estimate would directly affect the quality of the resulting motion reference, whose effect we discussed in Section 4.

Dependence on gentle depenetration in simulator. Our dynamics guidance assumes that the contact-rich policy remains close to the reference trajectory, since the derived force and torque targets are only valid when $q(t)$ is approximately preserved. Large depenetration velocity set in the simulators can induce large corrective impulses, which breaks our assumption and weaken the validity of the derived guidance.

Generalization beyond single motions. We currently train a separate policy for each motion. Extending the method to broader motion distributions is a natural next step, for example by distilling multiple motion-specific policies into a shared controller or by training a generative policy conditioned on contact geometries.

References

Raul Beltrán Beltrán, Julia Richter, Guido Köstermeyer, and Ulrich Heinkel. Climbing technique evaluation by means of skeleton video stream analysis. *Sensors*, 23(19):8216, 2023.

Rok Blagus, Bojan Leskošek, Luka Okršlar, Nace Vreček, and Tadej Debevec. Utility of video analysis and expert modelling for technique development in novice sport climbers: a randomized controlled study. *International Journal of Sports Science & Coaching*, 19(1):355–367, 2024.

Zhanxiang Cao, Yang Zhang, Buqing Nie, Huangxuan Lin, Haoyang Li, and Yue Gao. Learning motion skills with adaptive assistive curriculum force in humanoid robots. *arXiv preprint arXiv:2506.23125*, 2025.

Zhe Cao, Ilija Radosavovic, Angjoo Kanazawa, and Jitendra Malik. Reconstructing hand-object interactions in the wild. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12397–12406. IEEE, 2021.

Kyungsik Cha, Eun-Young Lee, Myeong-Hyeon Heo, Kyu-Cheol Shin, Jonghee Son, and Dongho Kim. Analysis of climbing postures and movements in sport climbing for realistic 3d climbing animations. *Procedia Engineering*, 112:52–57, 2015.

Aaron Ferguson, Ahmed AA Osman, Berta Bescos, Carsten Stoll, Chris Twigg, Christoph Lassner, David Otte, Eric Vignola, Fabian Prada, Federica Bogo, et al. Mhr: Momentum human rig. *arXiv preprint arXiv:2511.15586*, 2025.

Mohamed Hassan, Vasileios Choutas, Dimitrios Tzionas, and Michael J Black. Resolving 3d human pose ambiguities with 3d scene constraints. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 2282–2292, 2019.

Catalin Ionescu, Dragos Papava, Vlad Olaru, and Cristian Sminchisescu. Human3. 6m: Large scale datasets and predictive methods for 3d human sensing in natural environments. *IEEE transactions on pattern analysis and machine intelligence*, 36(7):1325–1339, 2013.

Angjoo Kanazawa, Jason Y Zhang, Panna Felsen, and Jitendra Malik. Learning 3d human dynamics from video. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5607–5616. IEEE, 2019.

Qiayuan Liao, Takara E. Truong, Xiaoyu Huang, Yuman Gao, Guy Tevet, Koushil Sreenath, and C. Karen Liu. BeyondMimic: From motion tracking to versatile humanoid control via guided diffusion. *arXiv preprint arXiv:2508.08241*, 2025.

Zhizheng Liu, Joe Lin, Wayne Wu, and Bolei Zhou. Joint optimization for 4d human-scene reconstruction in the wild. *The Fourteenth International Conference on Learning Representations*, 2026.

Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, 34 (6):248:1–248:16, October 2015.

Zhao Mandi, Yifan Hou, Dieter Fox, Yashraj Narang, Ajay Mandlekar, and Shuran Song. Dexmachina: Functional retargeting for bimanual dexterous manipulation. *arXiv preprint arXiv:2505.24853*, 2025.

Lea Müller, Hongsuk Choi, Anthony Zhang, Brent Yi, Jitendra Malik, and Angjoo Kanazawa. Reconstructing people, places, and cameras. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21948–21958. IEEE, 2025.

Alejandro Newell, Peiyun Hu, Lahav Lipson, Stephan R. Richter, and Vladlen Koltun. Comotion: Concurrent multi-person 3d motion. In *International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=qKu6KWPgxt.

Xue Bin Peng, Glen Berseth, and Michiel van de Panne. Terrain-adaptive locomotion skills using deep reinforcement learning. *ACM Trans. Graph.*, 35(4):81:1–81:12, July 2016. ISSN 0730-0301. doi: 10.1145/2897824.2925881. URL http://doi.acm.org/10.1145/2897824.2925881.

Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph.*, 37 (4):143:1–143:14, July 2018. ISSN 0730-0301. doi: 10.1145/3197517.3201311. URL http://doi.acm.org/10.1145/3197517.3201311.

Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Trans. Graph.*, 40(4), July 2021. doi: 10.1145/3450626.3459670. URL http://doi.acm.org/10.1145/3450626.3459670.

Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. Ase: Large-scale reusable adversarial skill embeddings for physically simulated characters. *ACM Transactions On Graphics (TOG)*, 41(4):1–17, 2022.

Soshi Shimada, Vladislav Golyanik, Weipeng Xu, and Christian Theobalt. Physcap: Physically plausible monocular 3d motion capture in real time. *ACM Transactions on Graphics (ToG)*, 39 (6):1–16, 2020.

Michiel Van de Panne and Alexis Lamouret. Guided optimization for balanced locomotion. In *Computer animation and simulation*, volume 95, pp. 165–177. Springer, 1995.

Michael Xu, Yi Shi, KangKang Yin, and Xue Bin Peng. Parc: Physics-based augmentation with reinforcement learning for character controllers. In *SIGGRAPH 2025 Conference Papers (SIGGRAPH '25 Conference Papers)*, 2025.

Xitong Yang, Devansh Kukreja, Don Pinkus, Taosha Fan, Jinhyung Park, Soyong Shin, Jinkun Cao, Jia-Wei Liu, Nicolás Ugrinovic, Anushka Sagar, et al. Sam 3d body: Robust full-body human mesh recovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7209–7219, 2026.

Wenhao Yu, Greg Turk, and C Karen Liu. Learning symmetric and low-energy locomotion. *ACM Transactions on Graphics (TOG)*, 37(4):1–12, 2018.

Ye Yuan and Kris M. Kitani. Residual force control for agile human behavior imitation and extended motion synthesis. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.

Ye Yuan, Shih-En Wei, Tomas Simon, Kris Kitani, and Jason Saragih. Simpoe: Simulated character control for 3d human pose estimation. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7155–7165. IEEE, 2021.

Haotian Zhang, Ye Yuan, Viktor Makoviychuk, Yunrong Guo, Sanja Fidler, Xue Bin Peng, and Kayvon Fatahalian. Learning physically simulated tennis skills from broadcast videos. *ACM Trans. Graph.*, 2023. doi: 10.1145/3592408.

Siwei Zhang, Yan Zhang, Federica Bogo, Marc Pollefeys, and Siyu Tang. Learning motion priors for 4d human body capture in 3d scenes. In *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 11323–11333. IEEE, 2021.

Siheng Zhao, Yanjie Ze, Yue Wang, C Karen Liu, Pieter Abbeel, Guanya Shi, and Rocky Duan. Resmimic: From general motion tracking to humanoid whole-body loco-manipulation via residual learning. *arXiv preprint arXiv:2510.05070*, 2025.

A Appendix

A.1 Details for kinematic reference extraction

We run an optimization with decision variable for body scale, root position and joint rotations, $x = [\beta,\ p,\ \theta]^T$; x_f is the decision variable for frame f. The objective function is

$$
\begin{aligned}
\min_{x} \quad & f(x) = \lambda(x) + \varphi(x) + \varsigma(x) + \eta(x) \\
\text{s.t.} \quad & v_{\text{ee}}(x_f) = 0, && \forall f,\ \forall \text{ee} \in \mathcal{S}(f), \\
& v^{\text{ang}}_{\text{parent(ee)}}(x_f) = 0, && \forall f,\ \forall \text{ee} \in \mathcal{S}(f), \\
& \text{SDF}(V_b(x_f)) \geq 0, && \forall f,\ \forall b \in \mathcal{B}_f
\end{aligned}
$$

where $\lambda(x) = \sum_f \sum_j \left\| P_j(x_f) - Y_{f,j} \right\|^2$ computes 2D projection cost across each frame f from $P_j(x_f)$, each joint j's forward kinematics position projected into the image space, and $Y_{f,j}$, the corresponding 2D joint positions from CoMotion; $\varphi(x) = \sum_f \sum_{i \in \mathcal{S}(f)} \left\| \text{FK}_i(x_f) - c_{f,i} \right\|^2$ is contact reach cost that pulls the world-space position of end effector i, $\text{FK}_i(x_f)$, towards its contact position prior $c_{f,i}$ for the end effectors $\mathcal{S}(f)$ that are in contact on frame f; $\varsigma(x) = \sum_f \left\| a_{r,f} \right\|^2$ computes temporal smoothness cost from the root acceleration in each frame $a_{r,f}$; $\eta(x)$ is a palm-orientation cost that softly aligns the palm normal with the contacted hold's surface normal at the contact point. The three constraints require the linear and angular velocities of contacting EEs to be zero, and the SDF of all body mesh vertices b, $\text{SDF}(V_b(x_f))$, to be positive. For readability we omit the per-term weights.

A.2 Derivation for equation 2

The equations of motion state

$$
M(q)\ddot{q} + h(q, \dot{q}) \;=\; S^\top \tau \;+\; B^\top F_{\text{ext}}, \tag{7}
$$

where F_{ext} can come either entirely from residual forces $F_{\text{res}} \in \mathbb{R}^6$ or entirely from contact forces $F_c \in \mathbb{R}^3$ per end effector. In the ideal result, F_{ext} comes entirely from contact forces,

$$
M(q)\ddot{q} + h(q, \dot{q}) \;=\; S^\top \tau \;+\; \sum_{i=1}^{N_{ee}} J_i(q)^\top F_{c,i}, \tag{8}
$$

where $F_{c,i}$ is the contact force on EE i. If, instead, we augment the policy's action with a residual force F_{res} applied at the pelvis and use it to fully drive the character, in a degenerate case where the policy makes no contact with the wall

$$
M(q)\ddot{q} + h(q, \dot{q}) \;=\; S^\top \tau' \;+\; B^\top F_{\text{res}}, \tag{9}
$$

where $F_{\text{res}} \in \mathbb{R}^6$ is the residual force applied at the root, and $B \in \mathbb{R}^{6 \times n}$ selects the six floating-base degrees of freedom of the character.

We partition the generalized coordinates as $q = (q_{\text{root}}, q_{\text{joint}})$ and split the Jacobian accordingly, $J_i = [\, J_i^{\text{root}} \;\; J_i^{\text{joint}} \,]$. Denote the six floating-base rows as $[\,\cdot\,]_{\text{root}}$ and the remaining actuated joint rows as $[\,\cdot\,]_{\text{joint}}$. Then equation 8 and equation 9 each separate into

$$
\left[M(q)\ddot{q} + h\right]_{\text{root}} = \sum\nolimits_i (J_i^{\text{root}})^\top F_{c,i}, \qquad \left[M(q)\ddot{q} + h\right]_{\text{joint}} = \tau + \sum\nolimits_i (J_i^{\text{joint}})^\top F_{c,i}, \tag{10}
$$

$$
\left[M(q)\ddot{q} + h\right]_{\text{root}} = F_{\text{res}}, \qquad \left[M(q)\ddot{q} + h\right]_{\text{joint}} = \tau', \tag{11}
$$

Require the two systems to produce the same motion $q(t)$, so their left-hand sides match exactly. Then the residual force that the degenerate policy applies at the root is exactly the net wrench that the ideal policy's end-effector contacts must supply:

$$
F_{\text{res}} \;\equiv\; \left[\, \sum\nolimits_{i=1}^{N_{ee}} J_i(q)^\top F_{c,i} \right]_{\text{root}}, \tag{12}
$$

and the joint torques required in the ideal policy τ are exactly defined by this equivalence and can be derived from the degenerate policy's joint torques τ':

$$
\tau = \tau' - \sum_i (J_i^{\text{joint}})^\top F_{c,i} \tag{13}
$$

In other words, we can solve for an estimate of $F_{c,i}$ from F_{res} to guide the policy's torque controls, as long as the poses $q(t)$ do not drift.