

KNOWING WHEN NOT TO HELP: Active Estimation of Human Reachability for Just-Right Robot Assistance

Ziang Liu^{1*}, Yunting Yan^{1*}, Christy Cheung^{1†}, Tailai Ying^{1†}, Bodong Liu^{1†}, Shiqin Tong¹,
Alexander Orkwis², Katherine Dimitropoulou³, Tapomayukh Bhattacharjee¹

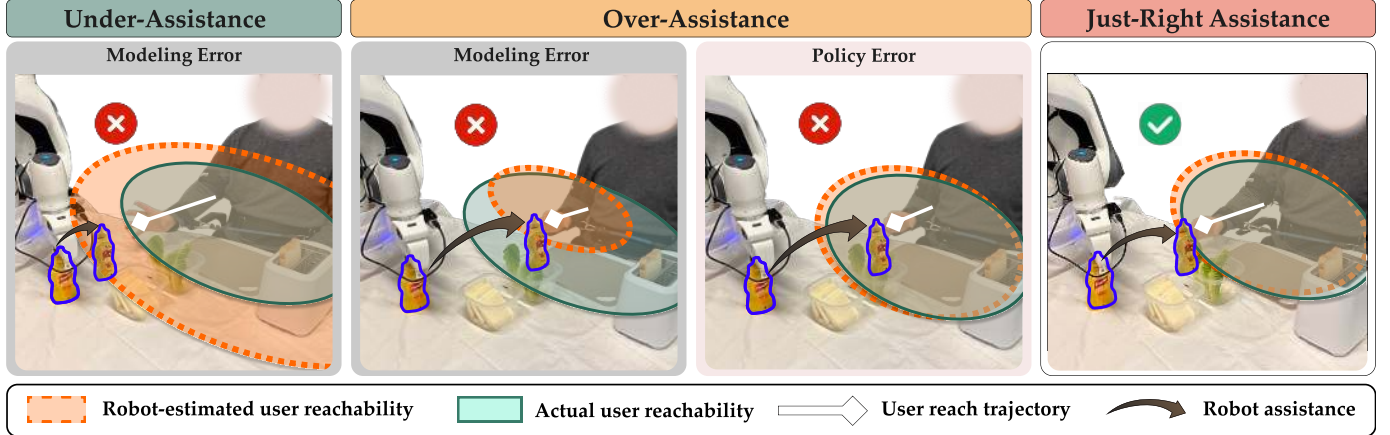


Fig. 1: We propose a framework that actively infers user-specific joint-space reachability from sparse interactions using a structured latent model, enabling robots to calibrate assistance and avoid under- or over-assistance during physical tasks.

Abstract—Robots that physically interact with humans must decide not only how and when to help, but also *when not to help*. In physical caregiving and collaborative manipulation, robots can over-assist by misestimating user capability or defaulting to helping when users can act independently. Physical functionality is highly individual, only partially observable, difficult to specify a priori, and assistance policies are often not grounded in user-specific ability, making calibrated intervention challenging.

We address this by actively inferring human joint-space reachability from sparse interaction. Our framework represents reachability using a compositional parametric model where a box constraint is deformed by local Gaussian primitives. We learn a latent space that decodes to these parameters and structure it using biomechanical anchors from musculoskeletal simulation and clinical anchors from retrieval-augmented reasoning over rehabilitation literature. The robot maintains a belief over this manifold and actively selects calibration queries to infer user-specific functionality.

We evaluate through computational experiments and real-robot studies with participants wearing resistance bands. Our method achieves ≈ 0.50 IoU within 20 queries. In sandwich making, reachability-aware assistance significantly improves user perception of physical engagement ($\chi^2(3) = 18.29, p < 0.001$) without increasing workload. In Action Research Arm Test (ARAT)-inspired manipulation, we demonstrate online adaptation. Website: <https://emprise.cs.cornell.edu/human-reachability/>.

¹Cornell University, Ithaca, NY. ²Binghamton University, Binghamton, NY. ³Columbia University, New York, NY.
{zl873, yy2244, csc263, tty6, bl576, tapomayukh}@cornell.edu;
sto4004@med.cornell.edu; aorkwis@binghamton.edu;
kd2524@cumc.columbia.edu

*Co-first authors. [†]Co-second authors.

I. INTRODUCTION

In physical robot caregiving, increased robot capability does not automatically lead to better assistance [2]. A robot that intervenes too frequently can suppress user agency [42] and contribute to “learned non-use” in rehabilitation, while a robot that intervenes too sparsely can cause task failure or injury. The robot often knows what the user wants to do (intent), but lacks an accurate estimate of whether the user can physically do it (functionality), or is simply ignorant of the user’s functionality. When this internal estimate is inaccurate, assistance becomes overly conservative or dangerously optimistic.

This paper focuses on a dominant failure mode in physical robot caregiving: *over-assistance caused by uncertainty about human functionality* [15]. Consider a robot assisting with sandwich making: it might retrieve an ingredient near the boundary of a user’s reachability, even if the user could safely have reached it. While technically completing the task, such assistance is unnecessary, intrusive, and agency-reducing. This failure stems not from poor control, but from an underestimate of reachability. Conversely, overestimating reachability risks unsafe under-assistance. Thus, assistance quality hinges on accurately estimating user reachability.

We model this through a *reachability envelope*, defined as the set of joint-space configurations a user can safely achieve. This envelope is highly user-specific, shaped by anatomy and impairment, but the robot must infer it from sparse, binary interaction outcomes. Existing assistive systems rely on

simplified geometric models [29], fixed reach limits [31], or offline motion capture and biomechanical analysis [39]. However, these models require difficult-to-tune a priori parameters and only provide point estimates, lacking explicit uncertainty quantification necessary for safe, conservative reasoning.

Estimating reachability online is fundamentally challenging. Exhaustively probing joint configurations is infeasible due to safety risks and the physical demands on the user. Direct self-report can describe broad limitations, but not the full joint-space boundary needed for assistance decisions. Reachability fluctuates due to short-term fatigue [35, 19] or long-term disease progression [24]. The robot must therefore infer a high-dimensional, structured reachability model from sparse observations, remain conservative under uncertainty, and adapt continuously as user reachability evolves. This leads to our research question: *How can we enable just-right robot assistance by efficiently inferring and maintaining an accurate estimate of a user’s reachability envelope from a small number of interaction queries?*

We address this using a model-based inference framework that treats reachability estimation as a Bayesian active learning problem [11]. We learn a structured latent representation of reachability from synthetic data, constrained by two forms of domain knowledge: biomechanical anchors from musculoskeletal simulation [13] and clinical anchors extracted from 270,000 rehabilitation articles via GraphRAG [43, 20]. During interaction, the robot maintains a belief over this latent manifold, actively selecting calibration queries to maximize information gain regarding the reachability boundary. This belief decodes into real-time reachability predictions to inform assistance decisions.

By explicitly modeling uncertainty in human reachability and updating reachability estimates in the loop, we design robot policies that refrain from unnecessary assistance when the user is capable, while remaining conservative when reachability is uncertain or degraded.

Contributions. This paper makes three contributions:

- A compositionally structured latent representation of human reachability, grounded in biomechanical and clinical anchors from musculoskeletal simulation and retrieval-augmented reasoning over rehabilitation literature, enabling interpretable reachability estimation without exhaustive per-user calibration.
- A Bayesian active learning framework for online, personalized reachability estimation, enabling sample-efficient adaptation through sparse user interaction.
- Real-robot validation with participants under emulated mobility constraints, demonstrating that reachability-aware assistance reduces over-assistance while maintaining high task success in manipulation and sandwich-making tasks.

II. RELATED WORK

A. Adaptive Assistance

Prior work in physical human–robot interaction highlights the importance of adaptive assistance grounded in explicit

modeling of human capabilities, shared autonomy, and environmental context to support users with mobility limitations [1, 7, 5, 38]. Existing approaches construct individualized capability representations, such as calibrated reachable workspaces or learned functional embeddings, to guide collaboration across users [12, 27, 3]. These offline calibration methods provide strong user-specific estimates when sufficient per-user data are available. Our goal is not to replace them, but to infer and adapt reachability in-situ from sparse interaction during assistance. Complementary efforts leverage multimodal sensing to estimate body geometry, contact forces, and task state for adaptive control [18, 30], or incorporate biomechanical models into robotic systems to enable capability-aware planning and assistive decision making [16, 25]. While task-space reachability models are useful for planning over end-effector workspaces, we model reachability in joint space because clinical range-of-motion limits, inter-joint coordination patterns, and biomechanical priors are naturally defined over joint configurations.

B. Shared Autonomy

Recent work in shared autonomy has shifted from fixed autonomy to explicitly modeling the fluid shift of control between humans and robots [1, 17, 2]. Complementary studies emphasize constraint-aware human intent modeling for physical collaboration, introducing tightly coupled estimation-and-control frameworks that account for both human operability and robot kinematic constraints [40]. In rehabilitation robotics, closely related assist-as-needed strategies provide assistance only when users deviate from desired behaviors, preserving autonomy and promoting active participation and engagement [45, 26]. Collectively, these works reflect a broader shift toward personalized shared autonomy that adapts robot assistance to individual human capabilities.

C. Active User Modeling

Active learning has been widely adopted in HRI to infer latent user attributes, such as preferences, skills, and reward functions, while minimizing interaction burden [4, 22]. Prior work has focused on efficiently estimating user preferences and task proficiency through informative queries, including dynamics-aware trajectory queries for reward inference [36], contrastive skill learning [8], causal-tree models for training difficulty estimation [9], and Bayesian formulations such as BKT-POMDP for multi-skill proficiency estimation in collaborative tasks [37]. We adopt active learning because reachability is unknown, user-specific, and defined over a high-dimensional joint space—making exhaustive or one-shot assessment costly. By selecting only the most informative queries near uncertain boundary regions, our approach enables sample-efficient online inference of user reachability.

D. Clinical Knowledge Grounding

Recent advances in clinical retrieval-augmented reasoning leverage curated medical knowledge to constrain inference




Fig. 2: Framework overview. The framework learns a structured latent model of joint-space reachability shaped by biomechanical and clinical anchors, then uses informative online calibration queries to infer user-specific functionality during interaction.

in safety-critical domains [46, 43], with extensions to robot-assisted procedures and embodied decision making [21, 28, 6]. We integrate these perspectives by using GraphRAG [10] to extract clinically grounded priors from large-scale rehabilitation literature and impose them as auxiliary constraints within an active learning framework for continuous joint-space human capability modeling. This enables data-efficient, individualized user modeling that is both interaction-aware and clinically informed, supporting adaptive assistive decision making.

III. PROBLEM FORMULATION

We formulate reachability estimation as a Bayesian active learning problem: given a limited query budget, sequentially select joint configurations to probe such that the robot can accurately infer the user’s reachability and make appropriate assistance decisions under uncertainty.

A. Reachability Representation

Let $\mathcal{Q} \subset \mathbb{R}^d$ denote the joint configuration space. A user’s reachability is characterized by a latent code $z_u \in \mathcal{Z}$, where $\mathcal{Z}$ is a learned latent space. This code parametrizes a continuous reachability field $f(\cdot; z_u) : \mathcal{Q} \rightarrow \mathbb{R}$, where $f(q; z_u) > 0$ indicates configuration q is reachable and $f(q; z_u) < 0$ indicates it is not. The reachability envelope is the zero level set:

$$\mathcal{R}_u = \{q \in \mathcal{Q} \mid f(q; z_u) \geq 0\}. \quad (1)$$

B. Observation Model

At step t , the robot queries configuration $q_t \in \mathcal{Q}$ as a target pose for the user to attempt, and observes a binary outcome

$y_t \in \{0, 1\}$ indicating whether the reach succeeded. We model observations as Bernoulli trials:

$$p(y_t = 1 \mid q_t, z_u) = \sigma\left(\frac{f(q_t; z_u)}{\gamma}\right), \quad (2)$$

where $\sigma(\cdot)$ is the logistic function and $\gamma > 0$ controls boundary softness.

C. Inference Objective

Given interaction history $\mathcal{D}_t = \{(q_i, y_i)\}_{i=1}^t$, the robot maintains and sequentially updates a posterior belief $p(z_u \mid \mathcal{D}_t)$ using the observation model. Under query budget T , the objective is to select queries maximizing information gain:

$$q_t^* = \operatorname{argmax}_{q \in \mathcal{Q}} \mathbb{I}(z_u; y \mid q, \mathcal{D}_{t-1}), \quad (3)$$

where $\mathbb{I}(\cdot; \cdot)$ denotes mutual information. This naturally prioritizes queries near the reachability boundary where uncertainty is highest.

For assistance decisions, the robot computes expected reachability $\bar{f}(q) = \mathbb{E}_{z \sim p(z \mid \mathcal{D}_t)}[f(q; z)]$ and its variance. High expected values indicate confident reachability (no assistance needed), low values indicate infeasibility (assistance required), and high variance indicates uncertainty in the reachability estimate.

IV. REPRESENTING USER REACHABILITY

We develop a parametric model of human joint-space reachability that enables efficient synthesis of diverse, clinically plausible impairment patterns. Instead of treating reachability as an unstructured geometric object, we model it compositionally: a baseline constraint captures independent joint




Fig. 3: Prior work models reachability using axis-aligned boxes or calibrated one-class SVMs [23, 27]. Boxes provide compact bounds but assume largely independent joint limits, while SVMs provide strong user-specific boundary estimates from calibrated data but yield data-dependent, variable-sized representations (Left). Our compositional model introduces additive and subtractive local Gaussian fields, enabling flexible data generation and augmentation (Middle). We validate representation fidelity on real user data and show scalability with additional fields (Right).

limits, while local primitives modify this baseline to represent coupled limitations. This structure preserves key topological properties observed in real human impairments while remaining compact enough to support sample-efficient inference.

We focus on four degrees of freedom in the right upper extremity: shoulder abduction-adduction, flexion-extension, internal-external rotation, and elbow flexion-extension. This abstraction captures gross motor reachability relevant to many functional tasks and focuses on joints where mobility limitations exhibit substantial inter-individual variation.

A. Compositional Parametric Model

Let $q \in \mathcal{Q} \subset \mathbb{R}^4$ denote a joint configuration. We represent reachability using a continuous field $f(q) \in \mathbb{R}$, where larger values indicate greater feasibility and the zero level set defines the reachability envelope.

Base kinematic constraint. We begin with a baseline that captures independent per-joint limits without coupling. For each joint i , lower and upper bounds $[l_i, u_i]$ define the feasible range, with asymmetric weights $w_i > 0$ allowing different penalty rates approaching the lower versus upper limits. The base field is:

$$d_{\text{base}}(q) = \text{smin}_{\alpha} \left(\left\{ \text{smin}_{\alpha}(q_i - l_i, w_i(u_i - q_i)) \right\}_{i=1}^4 \right), \quad (4)$$

where smin_{α} denotes the LogSumExp soft minimum with sharpness parameter α . This field assigns highest feasibility to mid-range configurations and decreases smoothly toward joint limits, treating all joints as independent.

Local coupling primitives. Joint limitations rarely occur independently. Motion in one joint often restricts or enables motion in others due to musculoskeletal geometry and neuromotor coordination patterns. We model these coupled constraints using local Gaussian primitives acting on joint-pair subspaces.

For joint pair (i, j) , each primitive k is defined by center $c_k \in \mathbb{R}^2$, anisotropic covariance Σ_k , amplitude a_k , and sign $s_k \in \{+1, -1\}$:

$$g_k(q_{(i,j)}) = a_k \exp \left(-(q_{(i,j)} - c_k)^{\top} \Sigma_k^{-1} (q_{(i,j)} - c_k) \right). \quad (5)$$

The final reachability field combines the baseline with these primitives:

$$f(q) = d_{\text{base}}(q) + \sum_{k=1}^K s_k \cdot g_k(q_{(i_k, j_k)}). \quad (6)$$

Primitives with negative signs carve out infeasible regions, modeling coupled restrictions such as reduced elbow extension during shoulder flexion. Primitives with positive signs locally expand the envelope, recovering feasible configurations that the axis-aligned box would otherwise exclude (Figure 3). We use $K = 30$ primitives (5 per joint pair: 3 negative, 2 positive).

Our compositional model provides an explicit parameterization supporting both fitting and generation. As shown in Figure 3, the model provides a compact parametric fit to observed reachability data while enabling synthesis of new reachability profiles through parameter perturbation. This generative capacity allows us to learn a latent space over diverse reachability patterns, as described in Section IV-B.

Model capacity validation. To validate model expressiveness, we fit the parametric structure to reachability data collected by Liu et al. [27], where participants wore braces inducing controlled mobility constraints. As shown in Fig. 3, models using both additive and subtractive primitives achieve lower reconstruction error than subtractive-only baselines. Error decreases monotonically with the number of primitives per joint pair. We use 5 primitives per pair as a practical trade-off between model expressiveness and compactness.

B. Latent Reachability Model

We learn a compact latent representation using an auto-decoder [33]. Each training sample is associated with a learn-




Fig. 4: Using a GraphRAG pipeline, the LLM retrieves user-specific clinical evidence and applies an OT-informed reasoning framework to estimate joint range-of-motion priors, while musculoskeletal simulation provides biomechanical anchors.

able latent code $z_u \in \mathbb{R}^{16}$, and a shared decoder network (3 residual blocks, 512 hidden units per block) maps z_u to the parameters of the compositional model:

$$\hat{\theta}_u = (\hat{l}_u, \hat{u}_u, \hat{w}_u, \{\hat{\beta}_{u,k}\}_{k=1}^{30}), \quad (7)$$

where $(\hat{l}_u, \hat{u}_u, \hat{w}_u)$ define the base box and $\hat{\beta}_{u,k}$ contains the primitive parameters (center, scale, amplitude, rotation) for each of the 30 Gaussian primitives. Ground-truth primitives are sorted by joint pair and center location to ensure consistent ordering. The fixed signs s_k are determined during data generation and remain frozen.

Training objective. We supervise the decoder using smooth ℓ_1 regression on the predicted parameters. Box parameters are trained in a center-span representation, and primitive parameters are regressed against canonical targets:

$$\mathcal{L}_{\text{recon}} = \mathcal{L}_{\text{box}} + \mathcal{L}_{\text{prim}}. \quad (8)$$

This reconstruction loss trains the decoder to accurately predict reachability parameters from latent codes.

C. Structuring the Latent Space

Parameter supervision alone does not guarantee that the latent space reflects meaningful structure. We impose additional regularization to encourage geometrically similar reachability envelopes to have nearby latent codes and to align latent directions with clinically interpretable patterns of impairment.

Geometric regularization. We preserve geometric similarity in the latent space by evaluating each training sample’s reachability field on a 16^4 grid to produce a binary occupancy signature. We compute pairwise Hamming distances $d_{\text{geo}}(i, j)$ between these signatures and train latent codes such that their Euclidean distances approximate these geometric distances:

$$\mathcal{L}_{\text{dist}} = \frac{1}{|\mathcal{P}|} \sum_{(i,j) \in \mathcal{P}} (\tilde{d}_z(i, j) - \tilde{d}_{\text{geo}}(i, j))^2, \quad (9)$$

where $\tilde{d}$ denotes mean-normalized distance over batch pairs $\mathcal{P}$. We add covariance regularization $\mathcal{L}_{\text{cov}}$ to prevent dimensional collapse and ℓ_2 regularization $\mathcal{L}_{\text{prior}}$ toward the origin.

Clinical anchors. We generate anchor samples representing clinical conditions identified by an occupational therapist as

clinically relevant and commonly observed in rehabilitation practice (e.g., adhesive capsulitis, stroke hemiparesis, cervical spinal cord injury). We construct these anchors using a GraphRAG pipeline over 270,000 rehabilitation articles (Fig. 4). We establish a keyword-based relational hierarchy to curate the literature database, parse articles to extract key propositions about conditions, interventions, outcomes, and ROM limitations, and organize them in a Neo4j [32] knowledge graph. Given a condition description, the system retrieves relevant propositions via graph traversal and applies chain-of-thought reasoning to infer ROM bounds $\mathbf{a}_k \in \mathbb{R}^8$ for each anchor k . A subset of these generated bounds was reviewed and verified by the occupational therapist to ensure clinical plausibility; additional conditions were generated by the model to provide broader coverage of the impairment space.

For each training sample, we decode ROM bounds $\hat{\mathbf{y}}_i$ via line search through the level set and compute soft assignment to clinical conditions c via Gaussian kernel:

$$p(c | i) = \frac{\sum_{k \in c} \exp(-\|\hat{\mathbf{y}}_i - \mathbf{a}_k\|^2 / (2\sigma^2))}{\sum_{k'} \exp(-\|\hat{\mathbf{y}}_i - \mathbf{a}_{k'}\|^2 / (2\sigma^2))}, \quad (10)$$

with $\sigma = 10$, where the sum in the numerator aggregates over all anchor variants belonging to condition c . We apply triplet loss with cosine distance and margin $m = 0.5$ to samples with confident condition assignments, pulling together latent codes with similar clinical profiles and pushing apart dissimilar ones.

Biomechanical anchors. Clinical diagnoses provide high-level categorization of impairment, but different conditions can manifest similar ROM limitations through distinct underlying mechanisms. To provide complementary low-level structure, we augment clinical anchors with biomechanical anchors derived from musculoskeletal simulation (Fig. 4, right). Using SCONe [13] and Hyfydy [14], we simulate anatomical movements under muscle weakness conditions to generate ROM bounds for specific muscle group deficiencies. Due to computational cost, we simulate single-degree-of-freedom movements starting from anatomical neutral to isolate individual joint limits. These anchors follow the same soft assignment and triplet loss formulation, grounding the latent

space in physiological mechanisms that may be shared across different diagnostic categories.

V. BAYESIAN ACTIVE LEARNING FOR REACHABILITY INFERENCE

Having learned a latent manifold of reachability patterns in Section IV, we now address the robot’s calibration problem: efficiently inferring which point on this manifold corresponds to a specific user. We treat this as a sequential inference problem where the robot selects joint configurations to query and updates its belief based on observed binary outcomes (feasible/infeasible). The inferred posterior enables the robot to predict reachability at unqueried configurations and modulate assistance decisions under uncertainty.

A. Posterior Representation

We represent the posterior as a categorical distribution over the training embeddings $\{z_1, \dots, z_N\}$ learned by the auto-decoder. Each embedding z_i decodes to a reachability field $f_i(q)$ via the fixed decoder network. We maintain unnormalized log-probabilities $\ell \in \mathbb{R}^N$, initialized to zero, which define the posterior via softmax.

This discrete parameterization bridges the generative model from Section IV and the sequential inference process: the categorical posterior naturally interpolates between training embeddings through weighted averaging, and the large training set made possible by our parametric reachability model provides broad coverage of clinically plausible impairment patterns.

B. Active Query Selection

Given this belief distribution, the robot must select queries that efficiently refine its estimate of user-specific reachability. At each round, we select the next query by maximizing predictive entropy over a candidate pool. For each candidate configuration q , we compute the model’s predicted probability that the user can reach q by averaging predictions across all training embeddings, weighted by their posterior probability:

$$\bar{p}(q) = \sum_{i=1}^N p(z_i | \mathcal{D}) \cdot \sigma\left(\frac{f_i(q)}{\gamma}\right), \quad (11)$$

where σ is the sigmoid function and γ controls boundary sharpness. The acquisition score is the binary entropy:

$$H(q) = -\bar{p}(q) \log \bar{p}(q) - (1 - \bar{p}(q)) \log(1 - \bar{p}(q)). \quad (12)$$

To focus queries near the reachability boundary, we compute the posterior-weighted mean field value $\bar{f}(q) = \sum_i p(z_i | \mathcal{D}) \cdot f_i(q)$ and gate the acquisition score: candidates with $|\bar{f}(q)| > \delta$ are excluded. This concentrates queries where the boundary is uncertain rather than in regions confidently feasible or infeasible.

We maximize entropy over a discrete set of candidate configurations sampled from the joint-space workspace. From the scored candidates, we select queries greedily subject to a minimum-distance constraint relative to previously queried configurations to prevent redundant sampling.

TABLE I: Ablation for parametric reachability prediction.

Metric	Ours	Ours (no reg)	Implicit
IoU $\uparrow$	0.971	0.982	0.827
# Islands $\downarrow$	1.32	1.31	32.23
Latent smoothness (ρ) $\uparrow$	0.998	0.294	0.844

TABLE II: Prior initialization for active learning.

Prior	IoU@0 $\uparrow$	IoU@10 $\uparrow$	IoU@20 $\uparrow$
Uniform	0.1140	0.5165	0.5353
Grid	0.5222	0.5788	0.6041
LLM	0.4706	0.6756	0.7988

C. Posterior Update and Prediction

Once a query is executed and the robot observes the outcome, the belief state is updated to incorporate this new evidence. After observing binary outcome $y \in \{0, 1\}$ at query q , we update the logits via Bayes’ rule. For each embedding z_i , we evaluate the likelihood:

$$p(y | q, z_i) = \begin{cases} \sigma(f_i(q)/\gamma) & \text{if } y = 1 \\ 1 - \sigma(f_i(q)/\gamma) & \text{if } y = 0 \end{cases} \quad (13)$$

We apply tempered likelihood updates to prevent premature posterior collapse:

$$\ell_i \leftarrow \ell_i + \eta \log p(y | q, z_i), \quad (14)$$

where $\eta < 1$ is the tempering parameter. For prediction, we compute the posterior-weighted average occupancy using Eq. (11) with the updated posterior weights. This Bayesian model averaging yields a calibrated probability reflecting both the model’s prediction and its uncertainty.

D. Assistance Modulation

Finally, the robot leverages this evolving posterior to make assistance decisions during task execution. For a goal configuration q_{goal} , we compute the posterior mean reachability via Eq. 11 with $q = q_{\text{goal}}$.

The assistance gain is:

$$\alpha(q_{\text{goal}}) = \sigma\left(-\frac{\bar{f}(q_{\text{goal}})}{\gamma_{\text{assist}}}\right). \quad (15)$$

When $\bar{f} \gg 0$, the configuration is predicted reachable and $\alpha \approx 0$ (minimal assistance). When $\bar{f} \ll 0$, the configuration is predicted infeasible and $\alpha \approx 1$ (full assistance). The transition sharpness is controlled by γ_{assist} . In this work, posterior uncertainty is used to guide query selection, while assistance modulation uses the posterior mean; extending the controller to use lower credible bounds or worst-case reachability estimates is left to future work.




Fig. 5: Left: reachability inference via active query selection. Right: active learning performance compared to baselines.

TABLE III: LLM joint limit prediction

Metric (MSE)	ShoulderF	ShoulderA	ShoulderR	ElbowF
Lower Bounds	29.06	11.77	12.12	27.18
Upper Bounds	24.94	20.93	10.08	25.84

VI. EXPERIMENTS

A. Parametric vs. Implicit Reachability Representation

We investigate whether explicit parameter prediction enables better-structured latent spaces than implicit geometry prediction. We compare our parametric decoder, which predicts box bounds and Gaussian primitive parameters that are then composed into a reachability field, against an implicit baseline that directly predicts signed distance fields using a SIREN network [41] with FiLM [34] conditioning (see Supplementary Material for architecture details). We also ablate our geometric regularization to isolate its contribution to latent space structure.

Metrics. We evaluate reconstruction IoU, topological consistency (mean number of disconnected components; ideal = 1.0), and latent smoothness (Spearman correlation between latent and geometric distances).

Key findings (Table I). Our parametric approach achieves superior reconstruction accuracy compared to the implicit baseline. More critically, the implicit baseline produces topologically invalid outputs: the mean decoded volume contains 32.23 disconnected components, violating the physical constraint that human reachability forms a single connected region. Our parametric approach produces near-ideal topology (mean 1.32 components) as a consequence of the compositional structure: the smooth Gaussian primitives modifying a connected baseline naturally preserve connectivity in most cases, though occasional disconnections can occur when primitives carve through the middle of the envelope.

Unregularized parametric prediction suffers from parameter-geometry ambiguity (different parameters can encode similar

shapes), causing manifold fragmentation and poor smoothness. While the implicit baseline (SIREN) achieves smoothness via continuous representation, our geometrically regularized approach matches this quality while preserving topological validity and reconstruction accuracy. This yields a structured latent space combining smooth geometry with physical correctness, essential for active learning to ensure posterior updates reflect meaningful, valid reachability changes.

B. Language-Based Prior Initialization

We investigate whether language models can reduce calibration effort by leveraging natural language descriptions of user conditions (e.g., “I have limited shoulder mobility due to a rotator cuff injury”). Using an LLM augmented with retrieval over a clinical knowledge base (Section IV-C), we map these descriptions to a latent prior via optimization-based inversion, enabling the robot to begin active learning from an informed starting point rather than a uniform prior.

Key findings (Table II). LLM-based initialization achieves competitive zero-shot performance (0.4706 IoU without physical queries) compared to 6 grid-sampled queries (0.5222 IoU), eliminating the burden of initial exploratory queries. As active learning proceeds, the LLM-initialized posterior converges faster: by iteration 10, it reaches 0.6756 IoU compared to 0.5788 (grid) and 0.5165 (uniform). The final performance gap persists at iteration 20 (0.7988 vs 0.6041 vs 0.5353), suggesting the LLM prior places the posterior in a favorable region of the latent space from which refinement is more effective. We validate LLM prediction accuracy on real stroke patient data (Table III), demonstrating that the approach produces clinically plausible ROM predictions from textual descriptions alone.

C. Reachability Inference via Active Query Selection

We evaluate our complete active inference approach on in-distribution and out-of-distribution users (Figure 5), measuring IoU over 100 queries averaged over 20 test users in each setting. Following prior calibrated reachability modeling, we use one-class SVM envelopes fit to the data from Liu et al. [27] as




Fig. 6: Robot Experiment Results: **(a)** Task conditions illustrating the effect of reachability model and policy on object placement: under-prediction places items too close (over-assistance), over-prediction places them too far (under-assistance), and our *just-right* policy targets the estimated reachability boundary; the *always-help* policy always repositions items near the user regardless of reachability. **(b)** Sandwich-making task procedure: participants retrieved four ingredients (bread, mustard, lettuce, cheese) sequentially. **(c)** Quantitative results across conditions: our *just-right* policy achieved the highest physical engagement and maintained full task success (4/4) while keeping physical workload moderate. **(d)** ARAT-inspired grasping task: the robot places a cube at the predicted reachability boundary on an ARAT shelf (red circle); if the user fails to grasp it, the robot issues active queries and repositions accordingly. **(e)** Reachability adaptation: the posterior over the user’s reachable workspace (dashed boundary) contracts as resistance band tension changed, and the robot updates object placement in response.

calibrated reference envelopes, and compare online inference methods by how efficiently they recover these envelopes from sparse queries. Thus, the SVM serves as the offline reference model rather than a competing online inference representation.

In-distribution performance. Our method achieves strong sample efficiency, reaching ≈ 0.50 IoU within 20 queries and converging to ≈ 0.65 IoU by 100 queries, tracking the oracle upper bound (≈ 0.70 IoU). Raw parameter space substantially underperforms (≈ 0.50 IoU), confirming that our latent representation provides essential structure. Space-filling fails to recover the boundary efficiently (≈ 0.10 IoU), reflecting the curse of dimensionality: uniform coverage rarely samples informative boundary regions in 4D joint space. Bayesian Active Learning by Disagreement (BALD) [22] performs comparably to our entropy-based acquisition (≈ 0.57 IoU), demonstrating that our latent space supports multiple acquisition strategies. Retrieval from real users plateaus at ≈ 0.35 IoU due to sparse coverage ($N=132$), while retrieval from generated users reaches ≈ 0.65 IoU, validating the value of parametric data augmentation and performing competitively with our probabilistic approach in this setting.

Out-of-distribution generalization. The oracle upper bound degrades to ≈ 0.62 IoU, reflecting reduced training

coverage. Our method maintains ≈ 0.57 IoU, demonstrating robust generalization. In contrast, retrieval methods collapse: nearest generated user drops to ≈ 0.48 IoU, while nearest real user falls to ≈ 0.30 IoU. Notably, our synthetic data generation provides substantial benefit over real data alone (0.48 vs 0.30), but probabilistic interpolation is essential to bridge the remaining gap ($0.48 \rightarrow 0.57$). When test users lie between training samples, retrieval-based matching cannot interpolate, whereas our categorical posterior smoothly synthesizes information across embeddings in the structured latent manifold. This performance gap under distribution shift demonstrates that posterior-based inference is critical for generalizable reachability estimation.

D. Real-Robot Evaluation: Knowing When Not to Help

We are interested in the extent to which our reachability model can predict accurate, personalized, and useful ROM in human-robot collaboration scenarios. We evaluate the system through a real-robot user study and an ARAT-inspired demonstration. The study protocol was reviewed and approved by the Institutional Review Board (IRB) of Cornell University.

1) *Multi-Step ADL Task (Sandwich Making)*: Meal preparation tasks such as sandwich making are fundamental ADLs that require coordinated actions across a distributed workspace

and can be challenging for individuals with upper-limb mobility limitations. We hypothesize (H_1) that *a reachability-aware assistance policy increases users' physical engagement compared to non-adaptive models and control policies that ignore individual reachability, without reducing task success.*

We evaluated this hypothesis through a real-robot user study using objective and subjective measures. The study employed a within-subject 3×2 factorial design with three reachability modeling conditions and two robot control policies: the *always-help* policy, which always repositioned ingredients near the user regardless of reachability, and the *just-right* policy, which intervened only when an item was predicted to be unreachable and placed it at the estimated reachability boundary to require final user effort. The three reachability modeling conditions were *under-predict* (underestimation/over-assistance, placing items too close), *over-predict* (overestimation/under-assistance, placing items too far), and our proposed method (*just-right*), as illustrated in Fig. 6(a). Since the *always-help* policy ignores reachability estimates entirely, it is invariant to the reachability modeling condition, collapsing the 3×2 design into four conditions: *under-predict*, *over-predict*, *always-help*, and *just-right*. Task success was recorded objectively, while physical engagement and physical workload were each assessed using 5-point Likert scales. For physical engagement, 1 indicated “not at all engaging” and 5 indicated “very engaging”; for physical workload, 1 indicated “not at all demanding” and 5 indicated “very demanding.”

The sandwich-making task was conducted using a Franka Emika Panda 7-DoF robotic arm. Participants assembled a sandwich using four ingredients distributed across a tabletop. Perception was provided by three Intel RealSense cameras: two external cameras placed in the room and one wrist camera mounted on the robot. SAM 3D Body [44] was used to estimate user joint configurations and task-space geometry, while the wrist-mounted camera enabled ingredient detection.

Ten student participants took part in the study. To simulate upper-limb mobility limitations, resistance bands were worn across all conditions, and condition order was counterbalanced across participants. An initial reachability estimate was obtained using 20 uniform-prior queries, after which the estimate remained fixed for the duration of the task. If a participant failed to retrieve a predicted-reachable item, they were encouraged to make additional attempts; after three failed attempts, the item was marked as unsuccessful. The average task completion time was 11 ± 3 min across all conditions, and we did not observe significant differences in task completion time between conditions. Inference ran in real time at 178 ± 1 Hz throughout.

Results are shown in Fig. 6. A Friedman test showed a significant effect of assistance condition on perceived physical engagement ($\chi^2(3) = 18.29$, $p < 0.001$). Our method yielded the highest engagement ($M = 3.6$, $Mdn = 4.0$), significantly outperforming both proactive assistance ($p = 0.010$) and over-assistance ($p = 0.006$), while maintaining full task success (4/4). Physical workload remained moderate ($M = 3.4$)

and was significantly lower than underestimation ($M = 4.4$, $p = 0.019$). These results demonstrate that explicitly targeting the reachability boundary enables the robot to support task completion while promoting meaningful physical participation.

2) *ARAT-Inspired Assistance Task*: This study used the same participant group as the sandwich-making task (Sec. VI-D1). Each participant experienced this task once, and the study focused exclusively on our proposed *just-right* policy. While the sandwich-making task evaluated assistance strategies under a fixed reachability estimate, this task was designed to demonstrate the framework’s ability to adapt *online* to changes in a user’s reachability over time. To simulate fatigue-induced reachability shifts, researchers increased resistance band tension during the task, requiring the robot to update its beliefs about the user’s reachability. A cube was placed on an ARAT shelf at the user’s estimated reachability boundary using the posterior from Sec. VI-D1. If the user failed to grasp the cube, the robot issued active queries to update the reachability posterior and repositioned the cube to the nearest predicted reachable location, demonstrating closed-loop adaptation as reachability changed. Example trials are shown in Fig. 6(d).

VII. DISCUSSION AND LIMITATIONS

Our framework demonstrates that explicitly modeling human capability through active inference allows robots to provide assistance that is both effective and preserves physical engagement. By grounding the latent space in biomechanical and clinical anchors, the system successfully bridges the gap between high-level diagnosis and low-level joint-space functionality. However, several limitations exist.

Kinematic abstraction: Our 4-DoF formulation captures gross arm reachability, but extension to higher-DoF arm–torso models still requires further attention.

Clinical population validity: Our user study recruited healthy participants with resistance bands to simulate upper-limb impairment. This protocol may not capture the full complexity of clinical conditions. Validation with clinical populations remains an important next step.

Query efficiency: Although our method achieves better performance and recovers from poor initialization substantially faster than uniform baselines, reducing this to single-digit queries remains future work. Future work should also explore integrating users’ mental states and fatigue signals into the active update cycle to maintain well-calibrated assistance over longer interaction durations.

ACKNOWLEDGMENTS

This work was partly funded by National Science Foundation IIS #2132846 and CAREER #2238792. We thank Xiaoman Yang for assistance with figure preparation and Malte Jung for logistical support during the robot experiments.

REFERENCES

- [1] Jimmy Baraglia, Maya Cakmak, Yukie Nagai, Rajesh Rao, and Minoru Asada. Initiative in robot assistance during collaborative task execution. In *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 67–74, Christchurch, New Zealand, March 2016. IEEE. ISBN 978-1-4673-8370-7. doi: 10.1109/HRI.2016.7451735.
- [2] Tapomayukh Bhattacharjee, Ethan K. Gordon, Rosario Scalise, Maria E. Cabrera, Anat Caspi, Maya Cakmak, and Siddhartha S. Srinivasa. Is More Autonomy Always Better? Exploring Preferences of Users with Mobility Impairments in Robot-assisted Feeding. In *2020 15th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 181–190, March 2020.
- [3] Andreea Bobu, Andi Peng, Pulkit Agrawal, Julie A Shah, and Anca D Dragan. Aligning human and robot representations. In *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, pages 42–54, 2024.
- [4] Maya Cakmak, Crystal Chao, and Andrea L. Thomaz. Designing Interactions for Robot Active Learners. *IEEE Transactions on Autonomous Mental Development*, 2(2): 108–118, June 2010. ISSN 1943-0612. doi: 10.1109/TAMD.2010.2051030.
- [5] Tom Carlson and Yiannis Demiris. Human-wheelchair collaboration through prediction of intention and adaptive assistance. In *2008 IEEE international conference on robotics and automation*, pages 3926–3931. IEEE, 2008.
- [6] Mingcong Chen, Siqi Fan, Guanglin Cao, Yun-hui Liu, and Hongbin Liu. Uspilot: An embodied robotic assistant ultrasound system with large language model enhanced graph planner. *arXiv preprint arXiv:2502.12498*, 2025.
- [7] Nazli Cila. Designing human-agent collaborations: Commitment, responsiveness, and support. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems*, pages 1–18, 2022.
- [8] Nathaniel Dennler, Stefanos Nikolaidis, and Maja Matarić. Contrastive learning from exploratory actions: Leveraging natural interactions for preference elicitation. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 778–788. IEEE, 2025.
- [9] Nathaniel Dennler, Zhonghao Shi, Uksang Yoo, Stefanos Nikolaidis, and Maja Matarić. Modeling personalized difficulty of rehabilitation exercises using causal trees. *arXiv preprint arXiv:2505.04583*, 2025.
- [10] Darren Edge, Ha Trinh, Newman Cheng, Joshua Bradley, Alex Chao, Apurva Mody, Steven Truitt, Dasha Metropolitansky, Robert Osazuwa Ness, and Jonathan Larson. From local to global: A graph rag approach to query-focused summarization, 2025. URL <https://arxiv.org/abs/2404.16130>.
- [11] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International conference on machine learning*, pages 1183–1192. PMLR, 2017.
- [12] Xiaofeng Gao, Luyao Yuan, Tianmin Shu, Hongjing Lu, and Song-Chun Zhu. Show Me What You Can Do: Capability Calibration on Reachable Workspace for Human-Robot Collaboration. *IEEE Robotics and Automation Letters*, 7(2):2644–2651, April 2022. ISSN 2377-3766, 2377-3774. doi: 10.1109/LRA.2022.3144779.
- [13] Thomas Geijtenbeek. Scone: Open source software for predictive simulation of biological motion. *Journal of Open Source Software*, 4(38):1421, 2019.
- [14] Thomas Geijtenbeek. The Hyfydy simulation software, 11 2021. URL <https://hyfydy.com>.
- [15] Matthew Gombolay, Reymundo Gutierrez, Shanellie Clarke, Giancarlo Sturla, and Julie Shah. Decision-making authority, team efficiency and human worker satisfaction in mixed human-robot teams. *Autonomous Robots*, 39, 07 2015. doi: 10.1007/s10514-015-9457-9.
- [16] Matthew Green, Yoon No Gregory Hong, Jinsook Roh, and Benjamin J Fregly. Computational modeling and simulation of closed chain arm-robot multibody dynamic systems in opensim. *Multibody System Dynamics*, 56(4): 313–334, 2022.
- [17] Raphaella Groten, Daniela Feth, Angelika Peer, and Martin Buss. Shared decision making in a collaborative task with reciprocal haptic feedback - an efficiency-analysis. In *2010 IEEE International Conference on Robotics and Automation*, pages 1834–1839, May 2010. doi: 10.1109/ROBOT.2010.5509906.
- [18] Yijun Gu and Yiannis Demiris. Vttb: A visuo-tactile learning approach for robot-assisted bed bathing. *IEEE Robotics and Automation Letters*, 9(6):5751–5758, 2024.
- [19] Dana Gutman, Samuel Olatunji, and Yael Edan. Evaluating levels of automation in human–robot collaboration at different workload levels. *Applied Sciences*, 11(16): 7340, 2021.
- [20] Haoyu Han, Yu Wang, Harry Shomer, Kai Guo, Jiayuan Ding, Yongjia Lei, Mahantesh Halappanavar, Ryan A Rossi, Subhabrata Mukherjee, Xianfeng Tang, et al. Retrieval-augmented generation with graphs (graphrag). *arXiv preprint arXiv:2501.00309*, 2024.
- [21] Pengfei Hao, Hongqiu Wang, Guang Yang, and Lei Zhu. Enhancing visual reasoning with llm-powered knowledge graphs for visual question localized-answering in robotic surgery. *IEEE Journal of Biomedical and Health Informatics*, 2025.
- [22] Neil Houlsby, Ferenc Huszár, Zoubin Ghahramani, and Máté Lengyel. Bayesian Active Learning for Classification and Preference Learning, December 2011.
- [23] Shafagh Keyvanian, Michelle J. Johnson, and Nadia Figueroa. Learning realistic joint space boundaries for range of motion analysis of healthy and impaired human arms. In *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*, pages 1–8, 2023. doi: 10.1109/Humanoids57100.2023.10375147.
- [24] Sumayya Khalid, Fady Alnajjar, Munkhjargal Gochoo,

- Abdulrahman Renawi, and Shingo Shimoda. Robotic assistive and rehabilitation devices leading to motor recovery in upper limb: a systematic review. *Disability and Rehabilitation: Assistive Technology*, 18(5):658–672, 2023.
- [25] Wajid Khalil Khalil. Opensim-based coupled lower limb rehabilitation robots. *Journal of Modern Rehabilitation*, 2025.
- [26] Maoqin Li, Jiaji Zhang, Guokun Zuo, Guang Feng, and Xueliang Zhang. Assist-as-needed control strategy of bilateral upper limb rehabilitation robot based on gmm. *Machines*, 10(2):76, 2022.
- [27] Ziang Liu, Yuanchen Ju, Yu Da, Tom Silver, Pranav N Thakkar, Jenna Li, Justin Guo, Katherine Dimitropoulou, and Tapomayukh Bhattacharjee. Grace: Generalizing robot-assisted caregiving with user functionality embeddings. In *2025 20th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 686–695. IEEE, 2025.
- [28] Chang Han Low, Ziyue Wang, Tianyi Zhang, Zhitao Zeng, Zhu Zhuo, Evangelos B Mazomenos, and Yueming Jin. Surgraw: Multi-agent workflow with chain-of-thought reasoning for surgical intelligence. *arXiv preprint arXiv:2503.10265*, 2025.
- [29] Kim Tien Ly, Valeriy Semenov, Mattia Risiglione, Wolfgang Merkt, and Ioannis Havoutis. R-Igp: A reachability-guided logic-geometric programming framework for optimal task and motion planning on mobile manipulators. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 14917–14923. IEEE, 2024.
- [30] Rishabh Madan, Skyler Valdez, David Kim, Sujie Fang, Luoyan Zhong, Diego T Virtue, and Tapomayukh Bhattacharjee. Rabbit: A robot-assisted bed bathing system with multimodal perception and integrated compliance. In *Proceedings of the 2024 ACM/IEEE international conference on human-robot interaction*, pages 472–481, 2024.
- [31] Veronique Maheu, Philippe S Archambault, Julie Frappier, and François Routhier. Evaluation of the jaco robotic arm: Clinico-economic study for powered wheelchair users with upper-extremity disabilities. In *2011 IEEE international conference on rehabilitation robotics*, pages 1–5. IEEE, 2011.
- [32] Neo4j. Neo4j - the world’s leading graph database, 2012. URL <http://neo4j.org/>.
- [33] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 165–174, Long Beach, CA, USA, June 2019. IEEE. ISBN 978-1-7281-3293-8. doi: 10.1109/CVPR.2019.00025.
- [34] Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron C. Courville. Film: Visual reasoning with a general conditioning layer. *CoRR*, abs/1709.07871, 2017. URL <http://arxiv.org/abs/1709.07871>.
- [35] Pradip Pramanick and Chayan Sarkar. Defatigue: Online non-intrusive fatigue detection by a robot co-worker. In *2018 27th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)*, pages 1129–1136. IEEE, 2018.
- [36] Dorsa Sadigh, Anca D. Dragan, S. Shankar Sastry, and Sanjit A. Seshia. Active preference-based learning of reward functions. In *Robotics: Science and Systems*, 2017.
- [37] Nicole Salomons, Emir Akdere, and Brian Scassellati. BKT-POMDP: Fast Action Selection for User Skill Modelling over Tasks with Multiple Skills. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*, pages 4243–4249, Montreal, Canada, August 2021. International Joint Conferences on Artificial Intelligence Organization. ISBN 978-0-9992411-9-6. doi: 10.24963/ijcai.2021/583.
- [38] Mario Selvaggio, Marco Cognetti, Stefanos Nikolaidis, Serena Ivaldi, and Bruno Siciliano. Autonomy in physical human-robot interaction: A brief survey. *IEEE Robotics and Automation Letters*, 6(4):7989–7996, 2021.
- [39] Ajay Seth, Jennifer L Hicks, Thomas K Uchida, Ayman Habib, Christopher L Dembia, James J Dunne, Carmichael F Ong, Matthew S DeMers, Apoorva Rajagopal, Matthew Millard, et al. Opensim: Simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement. *PLoS computational biology*, 14(7):e1006223, 2018.
- [40] Yifei Simon Shao, Tianyu Li, Shafagh Keyvanian, Pratik Chaudhari, Vijay Kumar, and Nadia Figueroa. Constraint-aware intent estimation for dynamic human-robot object co-manipulation. *arXiv preprint arXiv:2409.00215*, 2024.
- [41] Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, David B. Lindell, and Gordon Wetzstein. Implicit neural representations with periodic activation functions. *CoRR*, abs/2006.09661, 2020. URL <https://arxiv.org/abs/2006.09661>.
- [42] Eric T. Wolbrecht, Vicky Chan, Vu Le, Steven C. Cramer, David J. Reinkensmeyer, and James E. Bobrow. Real-time computer modeling of weakness following stroke optimizes robotic assistance for movement therapy. In *2007 3rd International IEEE/EMBS Conference on Neural Engineering*, pages 152–158, May 2007. doi: 10.1109/CNE.2007.369635.
- [43] Junde Wu, Jiayuan Zhu, Yunli Qi, Jingkun Chen, Min Xu, Filippo Menolascina, and Vicente Grau. Medical graph rag: Towards safe medical large language model via graph retrieval-augmented generation. *arXiv preprint arXiv:2408.04187*, 2024.
- [44] Xitong Yang, Devansh Kukreja, Don Pinkus, Anushka Sagar, Taosha Fan, Jinhyung Park, Soyong Shin, Jinkun Cao, Jiawei Liu, Nicolas Ugrinovic, Matt Feiszli, Jitendra Malik, Piotr Dollar, and Kris Kitani. Sam 3d body: Robust full-body human mesh recovery, 2026. URL

<https://arxiv.org/abs/2602.15989>.

- [45] Leigang Zhang, Shuai Guo, and Qing Sun. Development and assist-as-needed control of an end-effector upper limb rehabilitation robot. *Applied Sciences*, 10(19):6684, 2020.
- [46] Xuejiao Zhao, Siyan Liu, Su-Yin Yang, and Chunyan Miao. Medrag: Enhancing retrieval-augmented generation with knowledge graph-elicited reasoning for health-care copilot. In *Proceedings of the ACM on Web Conference 2025*, pages 4442–4457, 2025.