Unfolding the Literature: A Review of Robotic Cloth Manipulation

Alberta Longhini[1], Yufei Wang[2], Irene Garcia-Camacho[3], David Blanco-Mulero[3], Marco Moletta[1], Michael Welle[1], Guillem Alenyà[3], Hang Yin[4], Zackory Erickson[2], David Held[2], Júlia Borràs[3], Danica Kragic[1]

[1]Department of Robotics, Perception, and Learning, KTH Royal Institute of Technology, 100 44 Stockholm, Sweden; email: albertal@kth.se, marco.moletta@gmail.com, mwelle@kth.se, dani@kth.se
[2]The Robotics Institute, Carnegie Mellon University, Pittsburgh, USA, PA 15213; email: yufeiw2@andrew.cmu.edu, zerickso@andrew.cmu.edu , dheld@andrew.cmu.edu
[3]Institut de Robòtica i Informàtica Industrial, CSIC-UPC, Barcelona, Spain, 08028; email: igarcia@iri.upc.edu, dblancom@iri.upc.edu, galenya@iri.upc.edu, jborras@iri.upc.edu
[4]Department of Computer Science, University of Copenhagen, Copenhagen, Denmark, 2100; email: hayi@di.ku.dk

Xxxx. Xxx. Xxx. Xxx. YYYY. AA:1–31

https://doi.org/10.1146/((please add article doi))

Keywords

Textiles, Deformable Object Manipulation, Generalization, Physical Properties Variations, Tasks Variations

Abstract

The realm of textiles spans clothing, households, healthcare, sports, and industrial applications. The deformable nature of these objects poses unique challenges that prior work on rigid objects cannot fully address. The increasing interest within the community in textile perception and manipulation has led to new methods that aim to address challenges in modeling, perception, and control, resulting in significant progress. However, this progress is often tailored to one specific textile or a subcategory of these textiles. To understand what restricts these methods and hinders current approaches from generalizing to a broader range of real-world textiles, this review provides an overview of the field, focusing specifically on how and to what extent textile variations are addressed in modeling, perception, benchmarking, and manipulation of textiles. We finally conclude by identifying key open problems and outlining grand challenges that will drive future advancements in the field.

1. INTRODUCTION

Textile deformable objects such as clothing items or household objects like bed sheets and blankets are ubiquitous in our daily lives. Their usage spans applications from healthcare and domestic environments to the textile industry. Efforts to automate the manipulation and processing of these objects promise to enhance recycling and textile reuse while providing greater assistance to aging populations. Despite recent advances in manipulation tasks such as assistive dressing (1, 2), folding or bagging (3, 4), textile manipulation remains challenging as the deformable nature of these objects breaks fundamental assumptions in robotics such as rigidity, known dynamics models, and low dimensional state space (5). Specifically, when forces are applied to a deformable body, they not only move the object but also change its shape. From a physics point of view, understanding how the shape changes requires knowledge about the object's physical properties, such as stiffness or elasticity. From a perceptual point of view, properties such as shape, color, and material provide distinct signals to visual and tactile sensors. Endowing robots with skills to perceive, manipulate, and address the diversity of these textile properties presents a compelling avenue toward autonomous agents. However, current methods for textile manipulation tend to be tailored to specific objects that mirror the properties of the simulators used in their design (4, 6, 7, 8, 9).

This review consolidates recent methods and applications of deformable object manipulation, specifically textiles, highlighting the challenges associated with variations in their physical properties. We evaluate the progress made and identify key areas requiring further research to enhance the generalization and adaptive capabilities of robots in handling real-world textiles. Recent reviews have explored specific methods and applications related to this object category. Particular emphasis has been given to grasping (10) and caregiving scenarios (11, 12). In contrast, our review broadens the scope by analyzing generalization across different textile variations and applications. Building on the foundational work on perception, modeling, and manipulation of deformable objects (13, 14), our work offers a distinct perspective on enhancing the generalization and adaptability of perceptual and manipulation skills. Similarly to (5), we seek to identify ongoing challenges in modeling, perception, and control, further addressing the underexplored area of benchmarking.

The review is organized as follows. In Section 2, we provide the fundamentals about textiles, covering the variations of properties and tasks that will be discussed throughout the document. Section 3 reviews analytical and learning methods to model textile dynamics, highlighting their connection to textile physical properties. Section 4 identifies current approaches to perceive textile properties, whereas Section 5 covers current approaches for textile manipulations with a focus on techniques that enable generalization and adaptability to variations of properties and to what extent do current methods account for variations in physical and mechanical properties and tasks. In Section 6, we discuss currently available resources such as benchmarks and datasets, that enables evaluating this generalization, including benefits and limitations. We provide a thorough overview of application areas in Section 7. We close with a discussion about the interplay between modeling, perception, and manipulation, and future perspectives in Section 8.

2. FUNDAMENTALS

This section provides an overview of the fundamental aspects of textiles, detailing the definitions and characterizations of textile objects, exploring the variations in their physical

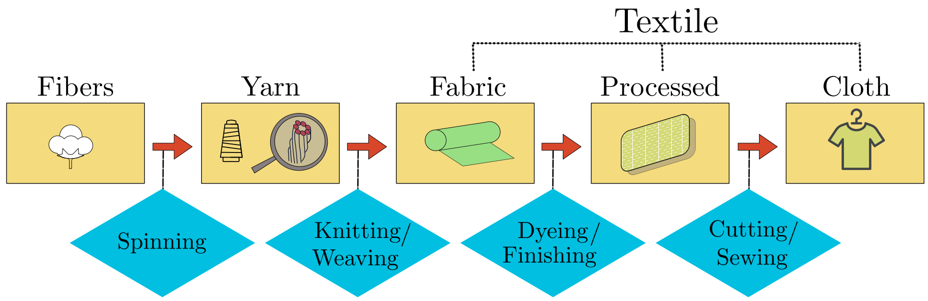

Figure 1

Manufacturing process of textiles: textile is an umbrella term that covers materials that are made of interlacing natural or synthetic fibers. The figure depicts the textile production process in its different stages, where yellow boxes represent materials and objects, while blue blocks specify the processing step.

properties, and examining the diversity in manipulation tasks. Understanding these foundational elements is essential for delineating the types of variations and challenges discussed in the subsequent sections on modeling, perception, and manipulation.

Textile: originates from the Latin *textilis*, meaning 'woven' and derived from the verb to weave.

2.1. Textiles, Fabrics, and Cloths

Textile has evolved from its initial reference, woven fabrics, to encompass a broad spectrum of objects. These include traditional woven fabrics, as well as deformable and flexible materials made from yarns or threads through various construction processes beyond weaving. Common terms used to describe textile objects are **fabric** and **cloth**. Despite being often overlapped in usage, they carry subtle distinctions tied to different phases of the production process, as shown in Fig. 1. Fabric is any thin, flexible material crafted from thread or yarn, fibers, polymeric film, foam, or their combinations, and is used in creating further products like clothing, requiring additional production steps (15). Cloth, while sometimes used interchangeably with fabric, typically refers to fabric that has undergone further processing or cutting. Everyday clothing items are predominantly made through weaving or knitting before being sewn, with high fashion exploring other methods more extensively.

The physical and mechanical properties of the final textile object are intrinsically tied to its manufacturing process. Yarns and threads, the building blocks of these textiles, are spun from various origins—animal, plant, mineral, synthetic, or blends thereof. Woven fabrics, produced by interlacing two sets of threads, are known for their hardness and non-elasticity, making them ideal for garments like shirts and jeans that retain creases. In contrast, knitted fabrics, created by interloping a single set of yarn, offer softness and elasticity in all directions and thus are used in wrinkle-resistant clothing such as t-shirts, which stretch uniformly in all directions to better conform to the body's shape. Additionally, further production steps involve cutting and coloring fabrics to create clothing with various shapes and colored textures. A comprehensive exploration of fabric properties is given in (16).

2.2. Variations of Object Properties

One of the major challenges in robotic manipulation is designing planning and control algorithms that generalize or adapt to variations that will be encountered in real-world deployments (17). In this review, we refer to **generalization** as the ability to apply the knowledge acquired from a specific set of textile properties to unseen variations of these properties. **Adaptation** is instead the ability to dynamically adjust strategies or models in response to changing conditions or novel variations in the object properties. Given the challenges of generalization and adaptation, it is crucial to first understand the variations in textile properties, which can be broadly classified into physical and mechanical properties. Fig. 2 shows an example of different variations of textile properties for four different objects.

2.2.1. Physical properties. The object's physical properties are inherent characteristics and can be observed and measured without subjecting the material to external forces or manipulations. These properties describe how a material appears under static or non-changing conditions. Among these properties, **Size** and **Shape** are crucial as they influence the robot's workspace, correlating with the object's function and category. The **Weight** influences how the textile deforms under gravity. For heavy and large textiles such as blankets, robots may need bimanual systems to manipulate them. The **Color** affects perception-based algorithms, essential for identifying and tracking the textile during manipulation tasks (1). **Fabric Material** composition, ranging from natural to synthetic fibers, dictates interaction forces between the textile and the robot, altering manipulation strategies (18). The **Construction Technique**, referring to the knitting or weaving processes, directly impacts a fabric's mechanical behavior by defining, for example, how tight or loose the construction pattern is, thus influencing the object elasticity and rigidity (19).

2.2.2. Mechanical properties. The mechanical properties are parameters that describe how a material responds to applied forces or manipulations. Cloth **Stiffness** or rigidity influences how it behaves under manipulation as it determines the resistance to deformation. Stiffness is a key determinant of a fabric's manipulation behavior, as it quantifies the textile's resistance to bending and deformation. **Elasticity**, or the capacity of a textile to stretch and recover to its original size after being deformed, is essential, particularly in applications like robotic dressing assistance (20). High elasticity can accommodate user movements and minimize safety risks during interaction. Elasticity and stiffness are properties of the material that are often characterized by Young's modulus and Poisson's ratio. The **Young's modulus** quantifies a material's ability to resist deformation under stress by showing the relation between stress and strain in the elastic region of the textile. The **Poisson's ratio** measures instead the ratio of transverse strain to axial strain in a stretched material. Finally, **Friction** refers to the resistance encountered when one surface slides over another. In the context of textiles, friction can vary significantly depending on the surface characteristics of the fabric (21). Friction properties affect how textiles move against surfaces, thus impacting how robots grasp and transport these materials and how they interact in contact with human skin or other clothing layers.

These mechanical properties vary based on the production process the textile undergoes. They can be measured through standardized methods from the textile industry (22) or through procedures tailored to robotics applications (23). Their interdependent influence in cloth deformation poses fundamental challenges to the analysis and understanding of how these properties affect robotic manipulations.

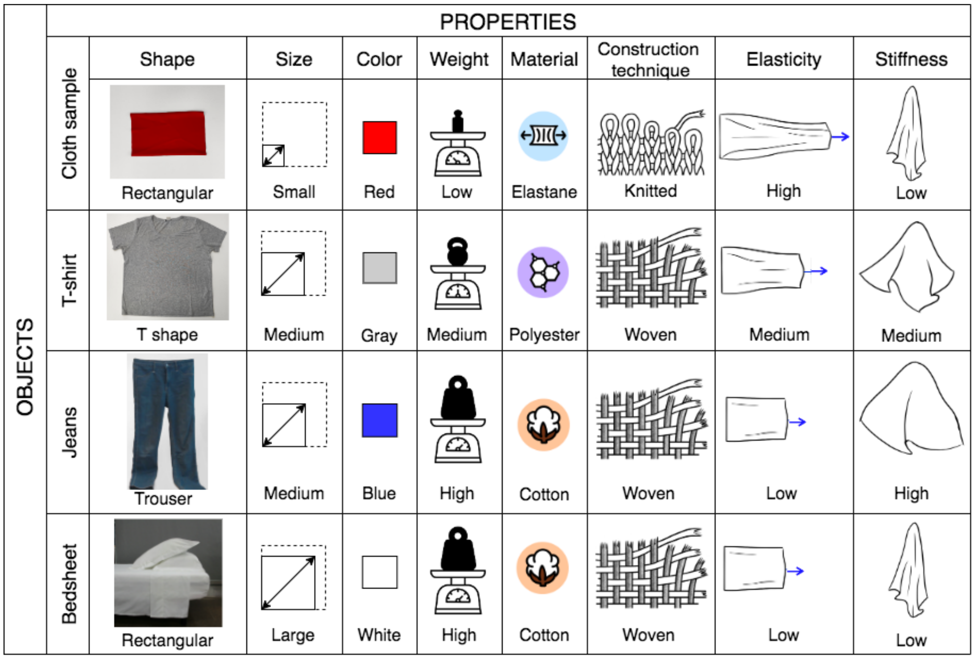

Figure 2

Object properties variations: Object differences rely on the variation of their physical and mechanical properties. The figure shows visual examples of how properties can vary between four objects from literature cloth sets. Measures can be found in (23).

2.3. Variations of Tasks

Manipulation of textiles spans a wide range of applications across domestic, healthcare, and industrial settings, including tasks such as laundry, tidying, dressing, sorting, and more. Each of these tasks presents unique challenges that necessitate the design and study of specific manipulation strategies that account for specific variations in the physical and mechanical properties of textiles discussed in the previous sections.

Following the taxonomy proposed by Mason (24), manipulation techniques can be characterized based on the nature of forces involved: kinematics, static forces, quasi-static forces, and acceleration forces. In the context of manipulating deformable objects, most current techniques can be broadly classified into two main categories (25): quasi-static manipulation, which involves slow motions allowing static equilibrium, and dynamic manipulation, which involves motions that include acceleration forces. Physical properties like material, as well as shape and size, significantly affect both types of manipulations. On the other hand, while mechanical properties play a fundamental role in tasks requiring dynamic manipulations due to the influence of acceleration forces, tasks involving quasi-static manipulations are less dependent on mechanical properties.

For instance, tasks such as classification or sorting highly depend on physical properties such as the shape of the cloth, as well as the material the textile is composed of (26). Similarly, tasks like flattening and folding rely heavily on understanding and manipulating the shape of the textile (27, 28, 29, 30). In these cases, elasticity does not directly impact the outcome of the manipulation. However, stiffness and friction still play a role in the deformed state in which the object is observed. Regarding shape properties, flattening and folding tasks are typically performed on textiles laid flat on a surface, simplifying the manipulation process by reducing the need to consider more complex topographical features such as loops, holes, or cylindrical parts. In contrast, tasks such as assistive dressing (31) or hanging an apron on a hook (32) require a nuanced understanding and control of the textile's topological features, including connectivity, holes, voids, and spatial relationships (33). These tasks are inherently more complex due to the need to manipulate the textile in three-dimensional space and adapt to its changing configuration. Mechanical properties, such as elasticity and stiffness, play a crucial role in these scenarios, as they influence how the textile drapes, stretches and recovers from deformation.

3. MODELING

Accurate modeling physical properties of textile objects is crucial for robotic manipulation, computer graphics, and material sciences. In this section, we introduce modeling techniques commonly used in robotics, focusing on how they characterize various properties and their effectiveness in capturing them. Rather than elaborating on the technical aspects of these techniques, which are extensively covered in other reviews, we aim to provide an overview of their utility and the types of physical and mechanical variations they address. The reality gap within commonly used simulators is also covered with a discussion about currently unmodeled phenomena and techniques for real-world alignment. For an in-depth technical discussion on textile modeling techniques, we refer readers to (13, 14, 34).

3.1. Physics-based Models

Physic-based models form the foundation for the most commonly used simulators in robotics. These models usually describe the cloth state through geometric representations such as particles or meshes. Physics-based simulation provides a reproducible and scalable mechanism to study physics-based cloth models and cloth dynamics across a range of environments and tasks. In what follows, we characterize different modeling techniques based on the representation used for the cloth.

Mesh-based models, such as mass-spring-damper models, represent the cloth as a mesh of interconnected masses and springs. The physical interaction between masses intuitively reflects the variation of textile stiffness and elasticity. These models are a popular choice implemented in many simulators, for example, MuJoCo (35) and SOFA (36). While these models are rather straightforward to implement, they are most suitable for simulating small deformation but not complex elastic effects with high fidelity. Additionally, tuning parameters for the springs and dampers to achieve realistic behavior can be relatively challenging, as these parameters do not directly correspond to physical meaning in the real cloth.

Particle-based methods in cloth simulation represent the material as a collection of discrete particles, each defined by properties like position, velocity, and mass. These particles are interconnected through holonomic constraints that describe their interactions. A widely used approach within this framework is Position-Based Dynamics (PBD), as implemented in simulation engines such as FLEX (37, 38). PBD focuses on satisfying constraints related to stiffness, elasticity, and collision by directly adjusting the positions of particles, with subsequent velocity updates derived from these positional changes. The advantages of PBD include rapid simulation speeds, enhanced stability, and the ability to simulate inextensible textiles and shear deformation. The ability to easily incorporate a variety of constraints makes it possible to model plastic deformations, as well as interactions with fluids and rigid bodies under flexible frictional models (37). While PBD offers visually plausible dynamics, it may not always align with physical realism. Moreover, the interpretation of simulation parameters into physical parameters, such as material modulus, remains challenging and often necessitates extensive tuning to achieve desired effects.

Continuous domain models leverage the principles of continuum mechanics rather than relying on discrete representations of particles or meshes. This approach treats physical quantities of objects as continuous fields, providing a more physically precise depiction of material deformations. The underlying partial-differential equations are commonly solved with the finite element method, which divides the field domain into small elements linked as meshes. While this approach can be computationally intensive, model parameters have a clearer physical interpretation for textile properties such as strain, stress and their relation through Young's modulus. Robotics-relevant simulators relying on this modeling technique are SOFA (36), Isaac Sim (39), and Bullet (40) but specifically for deformable objects. Recently, the Material Point Method (MPM) (41) emerges with a mix of particle representation, for which the particles carry material properties, while a spatial grid is used to compute forces and update the state of particles. By combining both representations, MPM is promising for rapid simulation of complex deformations observed in a cloth.

Besides forward prediction, **differentiable simulators** have recently gained popularity with efficient derivative evaluation through automatic-differentiation or adjoint methods (42). The feature extends existing model-based approaches for efficient policy learning and parameter identification, finding applications in cloth-related tasks (43, 44).

3.2. Data-Driven Models

Data-driven models have become a prominent alternative to traditional physics-based methods for simulating complex materials and deformations. Compared to physics-based techniques, data-driven models offer greater flexibility in defining the state space for modeling cloth dynamics, as well as lower computational complexity and easier parallelization.

These models learn cloth dynamics directly from data, which can range from handcrafted observations, such as cloth key points, to raw images and 3D point clouds or a latent representation of these. Image-based inputs have the benefit of not requiring knowledge of the 3D state of the cloth as the dynamics are learned either in the pixel space (30) or in a latent representation of the image (4). It also offers the advantage of working under partial observability (30). However, these often struggle with domain shifts that can occur with changes in camera positions, lighting, or background conditions, which impacts their generalization to new environments. Particle-based representations are often more robust to changes in visual conditions as they rely on 3D geometric representations such as particles or meshes rather than visual clues. These representations require specific architectures to efficiently capture local structures and handle data sparsity. Examples include PointNet++ (45) for unordered point sets and Graph Neural Networks (GNNs) (46) for mesh-based representations. Despite the advantage of particle-based representation in generalization, these models are generally more computationally intensive than their image-based counterparts. In addition, registering a mesh from partial depth data could also be non-trivial for textiles under large deformation.

The flexibility of data-driven models allows one to address variations of cloth properties by designing appropriate training schemes or conditioning the models on the object properties or a latent representation of these. Invariance to color changes is typically achieved during augmented training with domain randomization (47). In contrast, mechanical properties like stiffness and elasticity are commonly addressed through explicit or implicit conditioning. Explicit conditioning estimates these properties or their latent representations through perception and concatenates the estimation to the input model (19). Specifically for graph or mesh representations, properties like stiffness and elasticity can be induced as edge features in the mesh structure to bias model learning. Implicit conditioning, on the other hand, incorporates these properties by conditioning the model on a recent history of observations or employing techniques such as intuitive physics (48). These approaches allow models to adapt to changes in mechanical properties based on observed behaviors.

3.3. Reality Gap in Cloth Models

Different modeling techniques often need to account for a trade-off between computational cost and model accuracy. Trading accuracy for speed often leads to a gap between the dynamics in simulations and those in real-world scenarios. This gap between the model and how the physical cloth behaves in the real world is particularly relevant to real-world textile manipulation. Here, we identify gaps between the assumptions made by commonly used robot simulation and real textile properties, as well as progress to align simulation to the real world.

3.3.1. Unmodeled Phenomenas. Robot simulators generally favour fast and visually correct behaviours at the expense of physical realism. Geometrically, textiles are often modelled as 2D objects without thickness. This requires extra care when the task concerns stacking or folding textiles of multiple layers. The mechanical modelling, be it physics-based or data-driven, rarely goes down to the basic building blocks of fabrics as laid out in Section 2.1. The simulation hence neglects aspects like fabric weave, thread count, as well as the response to real-world factors such as humidity, temperature, or wear and tear. In computer graphics, yarn-level simulation (49, 50) has been studied to account for warp and weft interactions. However, this has not yet been integrated into robotics modeling pipelines. The simplification of simulating effects in smaller granularity also impacts the fidelity of how a cloth interacts with external objects. Textiles likely exhibit more complex sliding behaviours along different directions due to how they are constructed. This breaks the popular isotropic friction cone model in many simulators. Missing yarn-level simulation also influences modeling the tear-up of a cloth. Current robot simulators often assume invariant topology and uniformity about how a cloth is constructed. Graphics research has demonstrated the possibility of capturing these nuances with advanced simulation (49, 51). Furthermore, another phenomena that is rarely included in cloth models is the aerodynamics of the textiles. Since most garments are very light, even in the absence of wind, the air that surrounds them has a critical impact on highly dynamic cloth motions (52). Much research is still needed in order to have realistic, yet efficient physical models of cloth aerodynamics.

3.3.2. Aligning Simulation to the Real World. Physics simulation is essential for generating synthetic training data, exploring learned policies, and predicting the performance before real-world deployment. Thus, bridging the gap between the real world and simulation is crucial for generalizing manipulation skills to real-world textiles. A common approach to reduce this gap involves tuning the simulation parameters to align with reality. This process is often referred to as real2sim, or system identification, and can be classified into three categories (13): gradient-based techniques, global optimization techniques, and neural network-based techniques. For gradient-based optimization, differentiable simulators backpropagate the error between the real-cloth state and the simulated one for system identification. These methods, however, often face challenges with discontinuous loss landscapes, particularly in scenarios involving deformable objects. For scenarios where the analytical model is not differentiable, global optimization techniques such as Bayesian Optimization (BO) and Covariance Matrix Adaptation Evolution Strategies (CMA-ES) are often used without the need for gradients (53, 54). Global search methods can be computationally intensive and may struggle with scalability in high-dimensional spaces. A combination of the two has been proposed in (55), where they combine global search using BO with a semi-local search to retain the benefit of gradient-based optimization but integrate BO for the parts of the landscape that are intractable for gradient descent alone. Another approach is to combine Bayesian inference with neural networks to infer simulation parameters (56). One of the major benefits of this class of methods is capturing uncertainty in parameter estimates.

4. PROPERTIES PERCEPTION

The perceptual capabilities of robots encompass a variety of skills, including state estimation, segmentation, tracking, recognition, classification, and the identification of appropriate grasping points on cloth items. A comprehensive review of perception for grasping is presented in (11). Additionally, discussions on state estimation, parameter identification, and detection are provided in (14), where the focus was not only on textiles but on deformable objects in general. This section will specifically delve into the identification of textile properties, aiming to focus on perceptual capabilities that enable the generalization and adaptability of robots to variations in the textile properties introduced in prior sections.

Although robots in human environments constantly face uncertainty about object properties due to changing conditions, perceptual systems can reduce this uncertainty by providing sensory feedback (57). Perception can be passive, requiring no physical interaction with the objects of interest. Properties such as color, shape, and material can be estimated using passive perception (58). However, properties such as elasticity and friction cannot be observed in static scenarios. Therefore, interactive perception is necessary to obtain accurate estimates of these properties (59).

Designing the perception process relies on several key factors: the relevant properties for the tasks at hand, the available sensors, and the manipulator configuration. Figure 3 provides an overview of various design choices for estimating textile properties. This table also represents the literature discussed in the remainder of the section.

4.1. Physical Properties Perception

As noted in Section 2.2.1, physical properties like size, shape, weight, color, fabric material, and construction technique can often be observed and measured without external manipulation. Particularly, size and shape can be effectively inferred by observing textiles spread over a flat surface and utilizing contour or segmentation techniques (60). Data-driven tech-

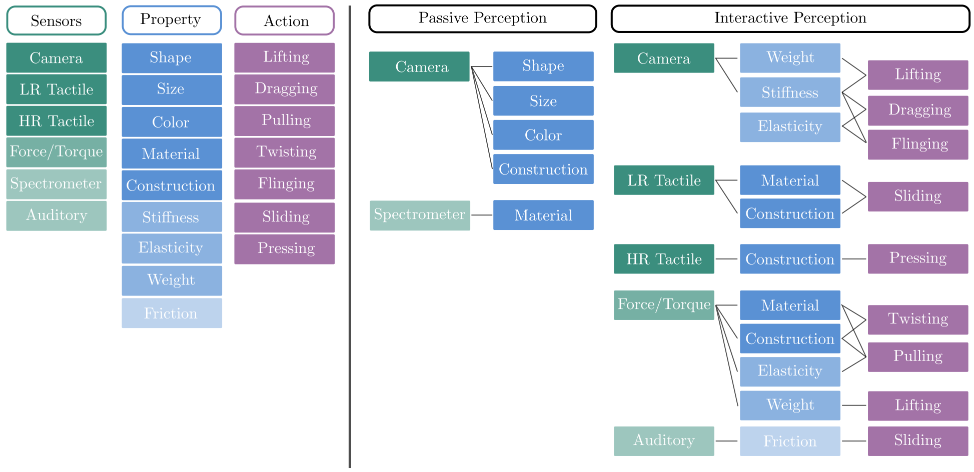

Figure 3

Overview of sensors, properties, and actions in passive and interactive perception for textile manipulation. On the left, the figure presents a list of sensors, properties, and actions. LR and HR for tactile sensors stand for low resolution and high resolution, respectively. The variation in opacity of the boxes reflects to what extent the sensor or the property has been explored for textile perception. On the right, the figure shows connections between sensors, properties, and actions based on the current literature. This side of the figure showcases how passive perception methods, using sensors like cameras and spectrometers, estimate properties that do not require physical interaction. In contrast, interactive perception integrates actions to improve the robot's understanding of mechanical properties, leveraging specific actions coupled with sensor feedback to accurately perceive properties such as weight, stiffness, and elasticity.

niques such as landmark detection or end-to-end classifiers have also been explored for shape classification (61, 62). However, textiles frequently arrive in deformed or crumpled states, complicating direct visual assessment. While some approaches have explored feature extraction processes to classify clothes from highly crumpled configurations (26), interactive perception is often necessary to enhance perceptual accuracy (63). Lifting interactions, for example, are commonly employed to infer the shape of the textile better and facilitate its classification by garment type (64, 65).

The estimation of cloth weight has been relatively underexplored within the robotics community. Pioneering investigations, primarily from the computer graphics domain, have demonstrated the potential of estimating the weight per area of textiles through video analysis of the object being influenced by external forces like wind (66, 67). This interaction could be analogously implemented in robotic systems as a flinging action (68), although this specific application has yet to be explored for weight estimation in a robotic context. Current approaches in the robotic community estimate cloths weights through the estimation of external forces/torques felt by the robot during motion tasks such as lifting (69). An alternative approach is to rely on a similarity network to infer cloth weight from lifting interactions using ground-truth cloth weight as supervision (63).

Textile material and construction techniques have received considerable attention in the robotic community to explore combinations of visual and haptic sensing specifically (70). Photometric stereo sensors can reconstruct the surface of a region of the cloth at the yarn level and exploit the 3D pattern to infer the material and construction properties of the object (71). Alternatively, high-resolution tactile sensing such as GelSight (72) has shown to be relevant for both material classification and construction technique classification of fabrics. These sensors provide high-resolution images, by pressing the sensor on the cloth (73). More specialized robot sensors, such as micro spectrometers, offer detailed insights into the yarn material composing the fabric (74).

Since both material and construction techniques influence the mechanical properties of cloth, interactive perception can provide additional insights to identify these properties accurately. Low-resolution tactile sensors like the BioTac, coupled with a contact microphone, have been explored to infer the material of textiles by exploring the feature of the signal recorded from a sliding interaction over the textile (70). Additionally, force-torque sensors used in conjunction with pulling and twisting actions have provided further insights into materials and construction techniques by observing how the textile responds to different mechanical stresses (18).

4.2. Mechanical Properties Perception

Mechanical properties of textiles, such as stiffness and elasticity, are critical in determining how a fabric will behave under various types of stress and strain. These properties are traditionally measured using sophisticated and costly systems that are standard in the textile industry. A prominent example is the Kawabata Evaluation System for Fabrics, which quantifies a fabric's response to controlled forces by measuring parameters such as stress-strain behavior at maximum load for the specific material being tested (75). Additionally, the computer graphics community has conducted in-depth studies of cloth elasticity using setups that combine pulling interactions with force and visual observations (76, 77). However, these setups are typically unsuitable for robotic applications due to their lack of real-time interaction and computational capabilities.

In the robotics community, stiffness and elasticity have been explored using a combination of force-torque sensors and camera observations during robotic pulling interactions (19). While this approach directly measures how textiles respond to controlled mechanical stresses, significant research focuses on methods using only camera observations to detect textile deformation under external forces like gravity or wind. Specifically, work using cameras can be broadly divided into two categories: 1) optimizing simulation parameters to reflect the behavior of real-world objects accurately and 2) applying external forces and observing resultant shape changes. The first category leverages 3D observation, such as point clouds, to optimize simulation parameters such as elasticity, bending, and stiffness (78, 79). The second category instead includes methods from the computer vision community, where 2D images are used to assess how a textile drapes under gravity or moves under external forces. These observations focus on visual features such as wrinkles or the dynamic motion of textiles over time (42, 80, 81). Notably, the studies of this last category do not involve robots, indicating a potential new research direction to replicate these techniques in robotic systems. Friction, on the other hand, remains largely underexplored. While some work has investigated haptic adjectives related to friction, direct estimation techniques are limited, with efforts mainly focusing on using low-resolution haptic and auditory sensors to infer friction-related features (82).

5. MANIPULATION

Robotic manipulation in human environments often confronts highly unstructured settings, necessitating the ability to manipulate objects with varied characteristics. In Section 2, we reviewed textile variations in terms of properties and tasks, and highlighted the complexity of addressing the variations of these real-world textiles. These complex environments underscore the importance of endowing robots with robust adaptation capabilities. Here, adaptation refers to the ability of the robot to adjust its strategies in response to changes in its operating environment or variations in the objects it manipulates. Effective manipulation skills require adaptability to both previously observed and novel variations of textile properties (57).

Successfully handling different clothing variations can be achieved through two primary mechanisms: generalization and adaptation (17). Generalization assumes that the distribution of knowledge acquired in the source domain will encompass the distribution of the target domain, resulting in a successful transfer. This assumption allows the robot to handle known variations, but it may not hold in the presence of novel variations typical of diverse human environments, where previously learned knowledge may fall short. A common framework to achieve generalization is domain randomization (47). In contrast, adaptation involves dynamically modifying the manipulation strategy or dynamics model by taking into account environmental changes. This adaptation can be achieved by parameterizing the dynamics model (48, 83) or the manipulation policy with the estimated cloth properties.

In the following sections, we group common manipulation strategies into three categories, provide a high-level overview of each, and discuss in more detail their adaptability towards different textile properties and task variations. A general overview of these methods is given in Table 1, including the tasks these are applied to as well as the variations of cloth physical and mechanical properties that are evaluated.

5.1. Model-based Manipulation

Model-based manipulation approaches refer to the class of methods that use a model of the textile to generate the manipulation strategy. The manipulation actions are usually obtained by planning with the model (54, 85, 86). These approaches can be categorized based on the methods used to construct the model, namely into analytical models and data-driven models.

5.1.1. Analytic Models. Analytic models, or physics-based models, build the textile model based on the laws of physics. Please refer to Section 3.1 for a more detailed discussion on the popular analytic models and the textile properties they can model. When using these models for manipulating a textile object, the initial step typically involves aligning the model dynamics to the real world to match the dynamics of the target textile (21, 103). Analytic models have the advantage of being applicable to different tasks as the model is usually task-agnostic, as long as the textile dynamics required by the task can be accurately modelled by the chosen analytic model.

There are two major drawbacks of using analytic models for manipulation. First, their generalization is limited by the accuracy of the physics model, which, as discussed in Section 3, is often only an approximation of real cloth dynamics (53). Second, each new object requires a separate system identification process, which can be tedious and difficult to scale for adapting to a great variability of textile mechanical properties.

Table 1 Overview of textile variations handled by methods in the literature. Methods that do not assess their generalization or adaptation over textile variations are marked with ✗, methods that address textile variations explicitly with ✓, and with (✓) when not explicitly mentioned but the evaluated textiles contains such variations.

Ref.	# O	Task	Physical Properties					Mechanical Properties			
			SH	SZ	CLR	M	CT	W	E	ST	F
Model-Based											
(4)	1	Fo	✗	✗	✗	✗	✗	✗	✗	✗	✗
(9)	1	Fo/Fl	✗	✗	✗	✗	✗	✗	✗	✗	✗
(48)	6	Fo	✗	✓	✓	✓	✓	✓	✓	✓	✗
(21)	7	Fo	✓	✓	✓	✗	(✓)	(✓)	✓	✓	✓
(30)	1	Fl	✗	✗	✗	✗	✗	✗	✗	✗	✗
(84)	3	Fl	✗	✗	✓	✗	✗	✗	✗	✗	✗
(85)	3	Fl	✓	✓	✓	(✓)	✗	(✓)	(✓)	(✓)	(✓)
(86)	5	Fl	✓	✓	✓	(✓)	✗	(✓)	(✓)	(✓)	(✓)
(54)	1	D	✗	✗	✗	✗	✗	✗	✗	✗	✗
Model-Free											
(7, 8, 28)	1	Fo	✗	✗	✗	✗	✗	✗	✗	✗	✗
(87)	2	Fo	✓	✓	✓	✗	✗	✗	✗	✗	✗
(88)	3	Fo	✓	✓	✓	✗	✗	✗	✗	✗	✗
(89)	3	Fo	✗	✗	✓	✓	(✓)	✓	(✓)	(✓)	✗
(90)	3	Fo	✓	✗	✓	✗	✗	✗	✗	✓	✗
(91)	5	Fo	✗	✗	✓	✗	✗	(✓)	✗	✗	✗
(92)	10	Fo	✗	✗	✓	✓	✗	✗	(✓)	(✓)	(✓)
(93)	10	Fo/Fl	✓	✓	✓	✗	✗	(✓)	✗	✗	✗
(94)	20	Fo	✓	✓	✓	✓	(✓)	(✓)	(✓)	(✓)	(✓)
(95)	1	Fl	✗	✗	✗	✗	✗	✗	✗	✗	✗
(47)	3	Fl	✗	✗	✓	✗	✗	✗	✗	✗	✗
(68)	3	Fl	✓	✓	✓	✗	✗	✗	✗	✗	✗
(96)	3	Fl	✓	✓	✓	✓	✗	✓	✓	✗	✗
(97, 98)	1	D	✗	✗	✗	✗	✗	✗	✗	✗	✗
(20)	2	D	✓	✓	✓	(✓)	✗	(✓)	(✓)	(✓)	(✓)
(99)	3	D	✗	✗	✓	✗	✗	✗	✗	✗	✗
(31)	5	D	✓	✓	✓	(✓)	✗	(✓)	(✓)	(✓)	(✓)
(100)	3	B	✗	✗	✓	(✓)	✗	✗	✗	✗	✗
Heuristic-Based											
(101)	4	Fo	✓	✓	✓	(✓)	✗	✗	✗	✗	✗
(29)	12	Fo	✓	✓	✓	(✓)	✗	(✓)	(✓)	(✓)	(✓)
(102)	25	Fo	✗	✓	✓	✓	✗	(✓)	(✓)	(✓)	(✓)

#O:Number of test objects, Fo:Folding, Fl:Flattening, D:Dressing, B:Bed Making, SH:Shape, SZ:Size, CLR:Color, M:Material, CT:Construction Technique, W:Weight, E:Elasiticy, ST:Stiffness, F:Friction.

5.1.2. Data-driven Models. Data-driven models learn from data about interactions with textiles. There can be many state representations for data-driven models as laid out in Section 3.2. The choice of the state representation and model architecture can greatly affect the generalization ability of the model towards different properties, such as the shape and texture of the textile. A particle-based state representation combined with graph neural networks (85, 86) has been shown to generalize better towards different shapes, geometries, and mechanical properties of the cloth compared to images (30) or latent-based state representations (4, 84), as the particle-based state representations align better with the underlying cloth physics. Using depth images (9) or point clouds (85, 86, 104) as the state representation also naturally makes the model invariant to the color and visual texture of the textiles. Due to the large amount of data needed for learning the model, most works use a simulator to generate the interaction data for learning the model (9, 30, 48, 84, 85, 86), with a few that does so in the real world (4).

As discussed in Section 3.2, data-driven models often address variations of cloth physical properties through domain randomization (30, 84, 85, 86). However, as current simulators (35, 37, 40) struggle in modeling a wide range of mechanical properties of the textiles, the generalization towards diverse mechanical properties that can be achieved through domain randomization remains limited. Still, recent work has shown that integrating a model conditioned on a latent representation of mechanical properties within a feedback-loop framework enables a model learned in simulation to adapt to textiles with diverse mechanical properties in the real world (48).

Most works in the literature learn an individual model for each type of manipulation task, e.g., assistive dressing (54), cloth folding (9, 48), cloth smoothing (9, 84, 86) or blanket covering (104). A few works have demonstrated that the learned model can be used to perform two tasks such as cloth folding and smoothing (9, 30, 85). Learning a single model that can generalize to multiple manipulation tasks remains underexplored.

5.2. Model-free Manipulation

Model-free manipulation approaches directly map the state or sensory observation of the textile to the manipulation action without having a model in the loop during inference. Compared to model-based approaches, they require no prior knowledge of the textile as no model needs to be constructed. There are two main approaches: reinforcement learning and imitation learning.

5.2.1. Reinforcement Learning. To apply reinforcement learning (RL) to a textile manipulation problem, the manipulation problem needs to be formulated as a Markov Decision Process (MDP) or a Partially Observable MDP (POMDP). The key in the formulation is the design of the state, action, and reward space of the MDP. As in the case of model-based methods, common choices for state representations of textiles include manually defined features such as key points on the cloth (92, 97), or results from a perception system such as an image (7, 47, 89) or point cloud (20, 31) of the textile. Common actions for manipulating textiles include action primitives such as pick-and-placing (47), dragging (90), flinging (68), or raw actions that control the delta movement or velocity of the robot end-effector (7, 31, 89). The reward function defines the desired outcome of the manipulation task based on the state of the environment.

Given the formulated MDP, reinforcement learning algorithms can be used to find a policy that maximizes the expected accumulated rewards, thus learning the necessary manipulation skills. Most works formulate the textile manipulation problem as a multi-step MDP (7, 31, 47, 89, 96) and learn a policy to maximize the accumulated reward, while some works (68, 90) formulate the problem into a bandit problem (an MDP with only 1 step), learn the reward function (usually in the form of a spatial-action value map (28, 68, 90)), and choose the best action as the one that maximizes the learned 1-step reward. Due to the large amount of data required by an RL algorithm, the need for a reward function, and the need to periodically reset the environment, such methods usually train a policy in a simulator and perform sim2real transfer (7, 31, 47, 89, 92, 96) with optional real-world fine-tuning (68). A few works (28, 87, 90) directly train in the real world, where the reward can be automatically computed from real-world perception systems (90), and the environment can be automatically (28, 90), or manually (87) reset.

Domain randomization is still the key technique for model-free RL methods to achieve generalization towards diverse textile mechanical and physical properties. In domain randomization, properties of the textile and environment are randomized during the training process, so the resultant policy generalizes to all varied properties. The randomized quantities can include the textile's shape (31, 68), size (7, 31), location and orientation (7, 31), texture and lighting (7, 47, 89), and mechanical properties (47, 89, 92). To achieve more informed domain randomization, contrastive learning can be used to compare pairs of real and simulated garment observations to learn a similarity metric (98), which is used to tune the simulation parameters to align the simulation and real-world garment observations. Again, one caveat of domain randomization is that the range of mechanical properties that can be randomized is limited by the fidelity of the simulator. There has been little work in the literature that explores addressing variations of textile properties via adaptation.

In terms of tasks, model-free RL methods have been applied to many textile manipulation problems including cloth smoothing and flattening (38, 47, 68, 90, 96), folding (7, 28, 87, 89, 90), placing (38), hanging (7, 32), blanket covering (105), and assistive dressing (20, 31, 97, 98). All these works learn an individual policy for each task; the goal of learning a single policy over multiple tasks or that can generalize to different manipulation tasks remains highly unexplored.

5.2.2. Imitation Learning. To apply imitation learning (IL) to textile manipulation problems, a dataset of expert demonstrations that solve the manipulation task needs to be first collected. Such demonstrations can be collected by a human (91, 94, 99) or using a scripted policy (3, 88, 95, 106). The most standard IL algorithm is behavioral cloning (91, 94, 95). Another approach is to directly map the actions in the demonstrations to the test object via learned correspondences (93).

The generalization ability of IL approaches highly depends on the amount of variations presented in the demonstrations. An extensive dataset that covers diverse objects with varied properties is essential for learning a policy that can generalize across different textiles. Many works that use IL for cloth manipulation show limited generalization towards different textile properties, demonstrating their method on a fixed textile (3, 8, 95, 106). Some works use depth images as observations, so the policy can be invariant to visual features such as color and texture (88, 91, 100). When the collected demonstrations are diverse enough, a stronger generalization can be achieved: Xue et al. (94) collected a dataset with hundreds of simulation shirts and 40 real-world shirts, and showed that the imitation policy can generalize to 20 real-world shirts with diverse shapes and materials. One caveat of IL methods is that collecting a sufficiently diverse and extensive set of demonstrations can be human labor-intensive.

Task-wise, IL methods have been applied to cloth smoothing (3, 93, 95), folding (8, 88, 91, 93, 94), twisting (106), bed-making (100), and dressing (99). Again, most works have only learned a single policy for one task (8, 91, 95, 106), with only a few that either learn a single backbone and different output branches on the same backbone (94), or a shared correspondence (93) for different tasks.

5.3. Heuristic-based Manipulation

The last discussed approach in textile manipulation is based on human-designed heuristic rules instead of learning the manipulation strategies from data. The heuristic rules vary depending on the target manipulation task. For example, for cloth smoothing, one commonly used rule is to detect the wrinkles of the cloth, and the manipulation action pulls the cloth in the perpendicular direction to the detected wrinkle direction (100). For grasping, one rule is to use the position and orientation of a wrinkle to compute the target location and orientation of the gripper (29). Another heuristic for a cloth article hanging in the air is to grasp key points that are identified based on the border geometry, e.g., corners for a towel or sleeves for a t-shirt (102), and use force sensors to trace the edge and get to the opposite corner (107). For folding, a heuristic folding motion can be achieved by leveraging gravity and moving the cloth such that the moved part is always vertical, which is called a "g-fold" (101). For assistive dressing, a common heuristic solution is to move the grasped garment along the forward direction of the human limb (54).

The generalization ability of such heuristically defined rules varies based on the target manipulation task and the assumptions made. Usually, these methods generalize well under the settings where the assumptions are satisfied and tend to have limited generalization when the assumptions are broken. Due to the diverse properties and configurations a textile can have, it can be hard to design a heuristic rule that would generalize to every situation one might encounter when manipulating textiles.

6. BENCHMARKS AND DATASETS

Benchmarking robotic manipulation plays a crucial role in understanding methods' capabilities and limitations, enabling standardized comparison. Particularly benchmarking of cloth manipulation is limited due to the absence of objective and consistent evaluation processes and the limited research on how different textiles influence the performance of a method.

In the following subsections, we will discuss the two main tools for benchmarking: datasets and benchmarks. Datasets can demonstrate the adaptability of a method to new data. However, when physical interaction is inherent to the task, the generalization of a method needs to be demonstrated through a benchmarked experimental validation, including standardization of the objects used.

6.1. Benchmarks

Most benchmarks tackling cloth objects evaluate manipulation in simulated environments, since they have the advantage of automatically generating various conditions to evaluate generalization. For instance, SoftGym (38) for deformable object manipulation and Assistive Gym (108) for assistive tasks, are simulation benchmarks for reinforcement learning that provides a set of simulated standardized environments. However, due to the sim2real gap, it is necessary to have benchmarks for evaluating real-world applications with physical objects. Garcia-Camacho et al. (109) proposes benchmarks with real executions for three cloth manipulation tasks including protocols, qualitative evaluation metrics, and several complexity levels based on the initial state of the cloth. Alternatively, Clark et al. (110) proposes four benchmark tasks for evaluating the performance of end-effectors in grasping clothing items, along with protocols to normalize crumpled configurations and metrics.

An effective way to measure the generalization of a method is through the use of a wide number of objects with varied properties. As it was done for rigid objects with the now widespread YCB object set (111), an extension focusing on textile objects is provided in (112), which was distributed among the participants of the cloth manipulation and perception competition (113). It includes a wide variety of cloth household objects with benchmarking guidelines for its use. However, an important issue in defining standardized textile object sets is stock continuity, preventing the maintenance of the same objects for extended time periods. To solve this issue, a method for building comparable textile object sets across different publications has been proposed in (23). It proposes a framework to characterize textile objects, enabling the quantification of variability on the textile properties listed in Section 2 of a given set of objects. With this characterization, the generalization of a method can be quantified based on the amount of variation that the objects offer. This idea goes in line with the one proposed for rigid objects in (114), where generalization is measured similarly but also adds other aspects of variability in background, table color, etc.

6.2. Datasets

Unlike benchmarks, datasets serve the purpose of providing data for designing or learning a task. The extent of generalization depends on the variation covered by the dataset. Datasets are often task-specific, ensuring that the data is relevant both for training a system and evaluating its performance. In cloth manipulation, the most common datasets include simulated 3D models (115, 116), RGB images (62, 117, 118) or depth (90, 119).

The largest existing datasets are for classic vision problems like cloth classification and landmark point detection, with datasets such as ClothesNet (115) in simulation or DeepFashion (62) with real images. Datasets for segmentation of people wearing clothes include CLOTH3D (116) in simulation or (120) with real images. Surface reconstruction is another classic perception problem applied to cloth (117). For cloth classification or landmark detection, images contain annotations of cloth type and location of landmarks, but usually, images come from the fashion industry, and so clothes are either flatted, hung on hangers, or worn by humans. For surface reconstruction, annotations need to have realistic images with the real mesh of the object, and therefore, existing datasets are all in simulation and subject to the sim2real gap discussed in Section 3.3.

For robotics, cloth classification and landmark point detection are also important, but they need to be identified during the stages of manipulation where clothes are in very complex configurations. Indeed, cloth classification is required but from crumpled states (26), or when grasped by one point (121). Other features that need to be identified are corners and edges (122) or wrinkles (123). The dataset (118) is an expansion from DeepFashion (62) to adapt it to robotic manipulation.

One of the challenges for datasets in cloth manipulation is labelling the ground truth of the deformation state in real images. Only a few datasets exist with depth or point clouds (119) and even less with labels of what points correspond to corners or edges during a manipulation (124). So far, to encode interactions, some datasets use RGB-D images where the action is annotated as a pick-up pixel point and a direction of motion in the image (90), with the limitation of only representing close-to-planar configurations. More complex actions appeared lately with Visual Language Models where a sequence of images is linked to a sequence of positions of the end-effector (125), with corresponding datasets.

Datasets rarely annotate the variability in the mechanical properties mentioned in Section 2, and some of the physical properties are covered depending on the requirements of the trained system. The main issue is the difficulty in labeling the deformation ground truth from images due to the severe self-occlusions, while datasets in simulation are only partly useful due to the sim2real gap.

7. APPLICATIONS AREAS

We outline the applications and challenges of manipulating deformable textile objects in scenarios requiring robots to adapt and generalize to their varying properties. Understanding the physical properties of textiles is crucial for a wide range of tasks in diverse sectors such as household chores, healthcare, and the textile industry. Table 2 provides a (non-exhaustive) overview of work categories in these sectors, organized by the frequency at which they are addressed in the literature. Tasks such as folding, smoothing, and dressing receive frequent attention from the community, whereas tasks like buttoning, dyeing, and washing remain rather underexplored. In what follows, we will discuss in detail these tasks and the requirements concerning variations of textile properties.

Table 2 Overview of Variation of Tasks addressed by the community and their frequency.

Frequency	**Household**	**Healthcare**	**Textile Industry**
Frequent (4+)	Folding (8, 29, 88, 90, 93, 126, 127) Smoothing (30, 47, 68, 93, 100) Ironing (128, 129, 130, 131, 132)	Dressing (31, 54, 127, 133, 134, 135, 136, 137, 138)	
Rare(2-3)	Hanging (7, 33) Sorting (26, 71) Wiping (139, 140)	Bedding (105, 126, 141) Bed-making (100, 126, 141) Bandaging (83, 142)	Recycling (71, 143)
Unaddressed (0-1)	Storing	Buttoning (144)	Manifacturing (145) Dyeing (146) Quality control Coloring Washing

7.1. Healthcare

As populations age worldwide (147), there is a growing opportunity for robotic systems that provide physical assistance with activities of daily living (ADLs) and healthcare tasks. Different assistive tasks involve manipulating textile objects, such as in robot-assisted dressing (127), bedding (141), bathing (138), and medical care (142). This section examines significant advancements in healthcare robotics, particularly robot-assisted dressing, bathing and bedding, bed-making, and medical care, highlighting opportunities and challenges in real-world applications.

Robot-assisted dressing is crucial for individuals with upper or lower extremity mobility impairments, requiring careful manipulation of deformable objects such as garments. This task necessitates the ability to adapt to various object variations while ensuring the safety and comfort of the patient. Research in robot-assisted dressing has led to the development of robots that can assist with putting on shirts (127, 133, 134), pants (135), and footwear (136), with challenges including ensuring physical safety, accurately modeling human-robot interactions during garment occlusions (54, 137, 138), and generalizing to different garments (31).

Bed-making and blanket manipulation represent significant opportunities for cloth handling in caregiving involving large, deformable textiles. Research has led to robotic systems capable of grasping and smoothing fitted sheets (141), folding and arranging blankets and towels (90, 126). Key to effective bed-making is leveraging physical properties like elasticity of bedsheets during robotic manipulation (100), also for tasks like autonomously covering and uncovering a person in bed (105).

In daily medical care, robotics research has introduced advances in handling soft materials such as gauze for bandaging (142) or adult diapers (148). These tasks, involving physical contact with the human body, underscore the importance of incorporating various sensory modalities and control techniques for effective manipulation of soft materials (83).

7.2. Household chores

Several instrumental activities of daily living (iADLs), such as laundry, cleaning with towels (fabric or paper), and hanging clothes, require dexterous textile manipulation. Cloth-like objects are ubiquitous in unstructured domestic environments and pose significant challenges to fully automating these activities. This section outlines methodologies and challenges in cloth sorting, smoothing, ironing, folding, hanging, and wiping tasks.

Cloth sorting involves categorizing garments and textiles by attributes such as item class (26), fabric type, construction, color, and quality, generally before washing or recycling. Accurate perception and classification of these variations can enhance sorting efficiency, enabling generalization across different textile batches.While recent methods have incorporated material identification through tactile feedback (71), few integrate physical interaction to accurately discern mechanical properties.

Robotic tasks such as smoothing, ironing, and wiping involve manipulating fabrics under varying physical conditions, which are primarily influenced by properties like friction and elasticity. Smoothing (30, 47, 68, 93, 100), typically performed before folding or wiping, necessitates to account for varying friction as fabrics transition from crumpled to flat states. Ironing further requires adjustments for temperature variations that affect the fabric's physical properties. Current methods mainly focus on ironing individual garments (128, 129, 130, 131, 132) and often fail to generalize across different fabric types.

Wiping involves tackling the challenges posed by variable surface friction, which can be influenced by the presence of dust or liquids on different surfaces (139, 140). Each task demands adaptive strategies to cope with the dynamic nature of fabric properties, highlighting the need for methods that can accommodate these variations, as discussed in Section 5.

Significant research in robotic manipulation of cloth has focused on robotic folding with practical applications in healthcare and domestic settings (8, 29, 88, 90, 93, 126, 127). However, challenges such as variability in garment properties (weight, friction, and shape) affect both quasi-static and dynamic manipulation (53), impacting generalization. Notably, material stiffness, as reflected in bending coefficients, significantly influences each fold, potentially accumulating errors and altering outcomes in the folding process (48).

Cloth hanging involves finding a stable configuration for a garment on a hanger (7) by identifying features like holes and loops (33). However, the impact of the physical properties of the garment on the deformation of these features remains largely underexplored.

7.3. Textile Industry

The textile industry, encompassing sectors such as fashion, automotive, and construction, presents numerous opportunities for autonomous robots with adaptive capabilities. This subsection focuses on specific applications such as manufacturing, dyeing, and recycling, directly linked to adaptability in handling, perceiving, and quality management of textiles. Given the extensive previous discussion on sorting, smoothing, ironing, and folding, this section will concentrate on these industry-specific applications.

In cloth manufacturing, robots are increasingly integrated into the cutting and sewing stages of garment production (145), with potential advancements enabling them to detect variations in fabric properties like thickness and stretchability. This ability could allow precisely adjusting techniques for each material, significantly reducing waste and enhancing resource efficiency, thus improving sustainability and garment quality.

The automation of dyeing processes significantly enhances sustainability by reducing dye and water usage (146). Currently, robots in dyeing processes are mainly used for loading and unloading yarn bobbins. Augmenting these robots to recognize variations in garment properties could optimize dye application, tailoring it to the specific needs of each garment.

Finally, recycling is crucial for sustainability (143). Similar to sorting, enhanced interactive perception of fabric types, construction methods, and material conditions can improve this process by accurately identifying textiles suitable for recycling (71).

8. DISCUSSION AND FUTURE PERSPECTIVES

In previous sections, we examined modeling, perception and manipulation separately. However, there is significant interplay among these domains, impacting generalization and adaptation capabilities. Offloading computational effort to modeling allows for model-based optimization of manipulation trajectories, reducing the burden on control and perception (21). However, real-world applications need perception techniques to align the parameters of the model to the real-world object, making perception crucial for generalization (48, 83). While model-free learning techniques like RL and IL learn end-to-end from raw data (89), reducing perception needs, they sacrifice sample efficiency and generalizability to novel variations of the environment and tasks. Nonetheless, not all components of manipulation tasks need end-to-end learning; perception modules can simplify sub-tasks like flattening wrinkles (27), grasp point detection (65), and folding plans, bypassing the need for extensive learning. As an example, perception can determine grasp points, how to reach the grasping point can be resolved with a standard planning algorithm, and how to optimize the manipulation once the cloth is grasped can be learned and performed by an end-to-end policy. Exploring when to switch between learning and heuristics is a promising research direction. This is particularly relevant with the advent of foundation models capable of reasoning about semantics, sequential tasks, and adapting to rules and human preferences.

In the remainder of the section, we further identify open problems and detail future research directions and grand challenges to foster the development of perception and manipulation skills that generalize to variations of textile properties and manipulation tasks.

8.1. Open Problems

To discuss generalization and adaptation in cloth manipulation, we reviewed fundamental textile properties and their complex, intertwined roles. However, the influence of each property in robotic manipulation is still underexplored, and the necessity of explicitly identifying each property remains unclear. An open avenue for research is identifying a subset of pertinent features for different manipulation tasks to determine the extent to which textile properties need to be identified.

Adding to the complexity of defining textile properties is their dynamic nature under varying conditions. Aging, wetness, dirt, and dryness can alter a textile's behavior, necessitating continuous sensing and adaptive perception in autonomous agents. Physical attributes like thickness, softness, and durability (73) are often evaluated alongside semantic descriptors like smooth, absorbent, hairy, and slippery (82). Advances in Large Language Models (LLMs) present potential to bridge these descriptors with a physical understanding of textile properties (149).

One key aspect for manipulating a variety of textiles and addressing the variation of properties is perception. While interactive perception with different types of sensors and exploratory actions can enhance adaptability and reduce uncertainty, it remains underexplored. Multimodal sensing emerges as a pivotal strategy in this context, integrating various sensory inputs like tactile, visual, and auditory data to provide a more holistic understanding of textile properties. This approach holds significant potential, as it allows robots to leverage multiple temporal information sources, compensating for the limitations of individual sensors.

A significant observation of our review of manipulation techniques is the limited amount of work demonstrating effective generalization to a wide range of deformable objects, with few exploring adaptation methods. Most efforts use domain randomization in simulations for sim2real transfer, often overlooking domain adaptation techniques that allow dynamic adjustment based on real-time feedback due to perceptual challenges. Similarly, state-of-the-art models, such as diffusion models, in combination with imitation learning techniques, remain underexplored. This leaves open questions about their applicability to deformable objects and the challenges that might arise due to the complexity of state estimation and high-dimensional state spaces of textiles.

Building on the necessity for improved generalization in manipulation techniques, benchmarking, standardization, and datasets emerge as essential yet challenging areas. The diversity in robotic embodiments, sensor configurations, and test sets complicates comprehensive comparisons. Developing universally applicable testing datasets faces challenges like stock availability and accurate property measurement. Although progress has been made, standardizing test sets and benchmarks remains a significant problem. Additionally, a major challenge in dataset availability and scalability is accurately labeling deformable object datasets, limiting their widespread distribution and application.

8.2. Grand Challenges

To push the boundaries of robotic manipulation of textiles, we identify the following critical challenges that may have major breakthroughs in the coming years:

(i) Perception of properties of novel objects for accurately estimating key physical and semantic properties of new textiles, enabling robots to reduce the uncertainties about the environment and adaptively handle a diverse range of textiles.

(ii) Adaptive multi-task and multi-modal agent capable of autonomously adapting to complex situations, performing long-horizon tasks, and using multi-modal sensory inputs to navigate uncertainties, enabling seamless integration into homes, industries, and healthcare facilities.

(iii) Novel datasets and benchmarks capturing real-world variations in object properties, physical interactions, and manipulation tasks to facilitate standardized benchmarking and enable consistent comparisons of robotic systems across research institutions.

The field of robotic manipulation of textiles is extensive and includes numerous important technological areas not mentioned here. Thus, the list provided is not exhaustive. The three grand challenges identified — perception of properties of novel objects, adaptive multi-task and multi-modal agents, and novel datasets and benchmarks — represent critical areas that have the potential to drive major advancements. These challenges encompass core perceptual technologies, general and adaptive capabilities, and standardized evaluation methods, aiming to enhance the flexibility and effectiveness of robotic systems in diverse real-world applications.

AUTHORS CONTRIBUTIONS

A.L. conceived the idea and the structure of the review, authored sections 1,2,4,8, coauthored sections 3,5, revised the full manuscript, created Figures 1 and 3, contributed to the creation of Tables 1 and 2, and coordinated the overall preparation and revision of the manuscript. Y.W. authored section 5, revised the full manuscript, and contributed to the creation of Table 1. I.G.-C. authored section 6, coauthored section 2, revised the full manuscript, and created Figure 2. D.B.-M. coauthored section 5, revised the full manuscript, and contributed to the creation of Table 1. M.M. authored section 7, revised the full manuscript, and contributed to the creation of Table 2. M.W. assisted with the delineation of the structure of the review, contributed to section 3, and revised the full manuscript. H.Y. coauthored section 3 and revised the full manuscript. Z.E. coauthored section 7 and revised the full manuscript. D.H. coauthored section 5 and revised the full manuscript. J.B. coauthored section 6 and revised the full manuscript. G.A. and D.K. revised the entire manuscript.

ACKNOWLEDGMENT

This work has been supported by the European Research Council (ERC-BIRD); the European Union's Horizon Europe Programme through projects SoftEnable (HORIZON-CL4-2021-DIGITAL-EMERGING-01-101070600) and IRE (HORIZON-CL4-2023-DIGITAL-EMERGING-01-101135082); the Swedish Research Council, the Knut and Alice Wallenberg Foundation, and the National Science Foundation under NSF CAREER grant number IIS-2046491; the project ROB-IN PLEC2021-007859 funded by MCIN/ AEI /10.13039/501100011033 and by the European Union NextGenerationEU/PRTR.

LITERATURE CITED

1. Kapusta A, Yu W, Bhattacharjee T, Liu CK, Turk G, Kemp CC. 2016. *Data-driven haptic perception for robot-assisted dressing*. In *25th IEEE International symposium on robot and human interactive communication*, pp. 451–58. Piscataway, NJ: IEEE
2. Chance G, Jevtić A, Caleb-Solly P, Dogramadzi S. 2017. A quantitative analysis of dressing dynamics for robotic dressing assistance. *Frontiers in Robotics and AI* 4:13
3. Seita D, Florence P, Tompson J, Coumans E, Sindhwani V, et al. 2021. *Learning to rearrange deformable cables, fabrics, and bags with goal-conditioned transporter networks*. In *IEEE International Conference on Robotics and Automation*, pp. 4568–75. Piscataway, NJ: IEEE
4. Lippi M, Poklukar P, Welle MC, Varava A, Yin H, et al. 2020. *Latent space roadmap for visual action planning of deformable and rigid object manipulation*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5619–26. Piscataway, NJ: IEEE
5. Zhu J, Cherubini A, Dune C, Navarro-Alarcon D, Alambeigi F, et al. 2022. Challenges and outlook in robotic manipulation of deformable objects. *IEEE Robotics Autom. Mag.* 29(3):67–77
6. Huang Z, Lin X, Held D. 2023. *Self-supervised cloth reconstruction via action-conditioned cloth tracking*. In *IEEE International Conference on Robotics and Automation*, pp. 7111–18. Piscataway, NJ: IEEE
7. Matas J, James S, Davison AJ. 2018. *Sim-to-Real Reinforcement Learning for Deformable Object Manipulation*. In *Proceedings of the 2nd Conference on Robot Learning*, ed. A Billard, A Dragan, J Peters, J Morimoto, pp. 734–743, vol. 87 of *Proceedings of Machine Learning Research*, pp. 734–743. N.p.: PMLR
8. Salhotra G, Liu ICA, Dominguez-Kuhne M, Sukhatme GS. 2022. Learning deformable object manipulation from expert demonstrations. *IEEE Robotics Autom. Lett.* 7(4):8775–82
9. Ma X, Hsu D, Lee WS. 2022. *Learning latent graph dynamics for visual manipulation of deformable objects*. In *2022 International Conference on Robotics and Automation*, pp. 8266–73. Piscataway, NJ: IEEE
10. Borras J, Alenya G, Torras C. 2020. A grasping-centered analysis for cloth manipulation. *IEEE Trans. Robotics* 36(3):924–36
11. Jiménez P, Torras C. 2020. Perception of cloth in assistive robotic manipulation tasks. *Natural Computing* 19(2):409–431
12. Wang L, Zhu J. 2023. Deformable object manipulation in caregiving scenarios: A review. *Machines* 11(11)
13. Arriola-Rios VE, Guler P, Ficuciello F, Kragic D, Siciliano B, Wyatt JL. 2020. Modeling of deformable objects for robotic manipulation: A tutorial and review. *Frontiers in Robotics and AI* 7:82
14. Yin H, Varava A, Kragic D. 2021. Modeling, learning, perception, and control methods for deformable object manipulation. *Sci. Robotics* 6(54):eabd8803
15. Kadolph SJ. 2007. *Textiles*. Upper Saddle River, N.J.: Pearson Prentice Hall, 10th ed.
16. Grishanov S. 2011. Structure and properties of textile materials. In *Handbook of textile and*

industrial dyeing. Elsevier

17. Cui J, Trinkle J. 2021. Toward next-generation learned robot manipulation. *Sci. Robotics* 6(54):eabd9461
18. Longhini A, Welle MC, Mitsioni I, Kragic D. 2021. *Textile taxonomy and classification using pulling and twisting*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 7564–71. Piscataway, NJ: IEEE
19. Longhini A, Moletta M, Reichlin A, Welle MC, Kravberg A, et al. 2023. *Elastic context: Encoding elasticity for data-driven models of textiles*. In *IEEE International Conference on Robotics and Automation*, pp. 1764–70. Piscataway, NJ: IEEE
20. Sun Z, Wang Y, Held D, Erickson Z. 2024. Force-constrained visual policy: Safe robot-assisted dressing via multi-modal sensing. *IEEE Robotics Autom. Lett.* 9(5):4178–4185
21. Li Y, Yue Y, Xu D, Grinspun E, Allen PK. 2015. *Folding deformable objects using predictive simulation and trajectory optimization*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 6000–06. Piscataway, NJ: IEEE
22. Cusick G. 1968. The measurement of fabric drape. *The Journal of The Textile Institute* 59(6):253–60
23. Garcia-Camacho I, Loghnini A, Welle MC, Alenyà G, Kragic D, Borras J. 2024. *Standardization of cloth objects and its relevance in robotic manipulation*. In *IEEE International Conference on Robotics and Automation*. Piscataway, NJ: IEEE
24. Mason MT. 2001. *Mechanics of robotic manipulation*. MIT press
25. Blanco-Mulero D. 2024. Towards efficient robotic manipulation of deformable objects by learning dynamics models and adaptive policies
26. Sun L, Aragon-Camarasa G, Rogers S, Stolkin R, Siebert JP. 2017. *Single-shot clothing category recognition in free-configurations with application to autonomous clothes sorting*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 6699–6706. Piscataway, NJ: IEEE
27. Sun L, Aragon-Camarasa G, Cockshott P, Rogers S, Siebert JP. 2014. *A heuristic-based approach for flattening wrinkled clothes*. In *Towards Autonomous Robotic Systems*, ed. A Natraj, S Cameron, C Melhuish, M Witkowski, pp. 148–160, pp. 148–160. Berlin, Heidelberg: Springer
28. Lee R, Ward D, Dasagi V, Cosgun A, Leitner J, Corke P. 2021. *Learning arbitrary-goal fabric folding with one hour of real robot experience*. In *Proceedings of the 4th Conference on Robot Learning*, ed. J Kober, F Ramos, C Tomlin, pp. 2317–27, vol. 155 of *Proceedings of Machine Learning Research*, pp. 2317–27. N.p: PMLR
29. Doumanoglou A, Stria J, Peleka G, Mariolis I, Petrik V, et al. 2016. Folding clothes autonomously: a complete pipeline. *IEEE Trans. Robotics* 32(6):1461–78
30. Hoque R, Seita D, Balakrishna A, Ganapathi A, Tanwani AK, et al. 2022. Visuospatial foresight for physical sequential fabric manipulation. *Auton. Robots* 46(1):175–99
31. Wang Y, Sun Z, Erickson Z, Held D. 2023. *One policy to dress them all: Learning to dress people with diverse poses and garments*. In *Robotics: Science and Systems XIX*, ed. K Bekris, K Hauser, S Herbert, J Yu, no. pap. 008. N.p.: Robotics: Science and Systems Foundation
32. Antonova R, Shi P, Yin H, Weng Z, Jensfelt DK. 2021. *Dynamic environments with deformable objects*. In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, ed. J Vanschoren, S Yeung, vol. 1
33. Antonova R, Varava A, Shi P, Carvalho JF, Kragic D. 2021. *Sequential Topological Representations for Predictive Models of Deformable Objects*. In *Proceedings of the 3rd Conference on Learning for Dynamics and Control*, ed. A Jadbabaie, J Lygeros, GJ Pappas, PA Parrilo, B Recht, CJ Tomlin, MN Zeilinger, pp. 348–360, vol. 144, pp. 348–360. N.p.: PMLR
34. Hou YC, Sahari KSM, How DNT. 2019. A review on modeling of flexible deformable object for dexterous robotic manipulation. *International Journal of Advanced Robotic Systems* 16(3)
35. Todorov E, Erez T, Tassa Y. 2012. *MuJoCo: a physics engine for model-based control*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 5026–33. Piscat-

away, NJ: IEEE
36. Faure F, Duriez C, Delingette H, Allard J, Gilles B, et al. 2012. Sofa: a multi-model framework for interactive physical simulation. *Soft Tissue Biomech. Model. for Comput. Assist. Surg.* :283–321
37. Macklin M, Müller M, Chentanez N, Kim TY. 2014. Unified particle physics for real-time applications. *ACM Trans. Graph. (TOG)* 33(4):1–12
38. Lin X, Wang Y, Olkin J, Held D. 2020. *SoftGym: benchmarking deep reinforcement learning for deformable object manipulation.* In *Proceedings of the 4th Conference on Robot Learning*, ed. J Kober, F Ramos, CJ Tomlin, pp. 432–48, vol. 155 of *Proceedings of Machine Learning Research*, pp. 432–48. N.p.: PMLR
39. Macklin M, Erleben K, Müller M, Chentanez N, Jeschke S, Makoviychuk V. 2019. Non-smooth newton methods for deformable multi-body dynamics. *ACM Trans. Graph.* 38(5)
40. Coumans E, Bai Y. 2016–2021. Pybullet, a python module for physics simulation for games, robotics and machine learning. `http://pybullet.org`
41. Jiang C, Schroeder C, Teran J, Stomakhin A, Selle A. 2016. *The material point method for simulating continuum materials.* In *ACM SIGGRAPH 2016 Courses*, SIGGRAPH '16. New York, NY, USA: Assoc. Comput. Mach.
42. Larionov E, Eckert ML, Wolff K, Stuyck T. 2022. Estimating cloth elasticity parameters using position-based simulation of compliant constrained dynamics. *arXiv preprint arXiv:2212.08790*
43. Qiao YL, Liang J, Koltun V, Lin MC. 2020. *Scalable differentiable physics for learning and control.* In *Proceedings of the 37th International Conference on Machine Learning.* N.p.: JMLR.org
44. Chen S, Xu Y, Yu C, Li L, Ma X, et al. 2023. *DaxBench: Benchmarking Deformable Object Manipulation with Differentiable Physics.* In *The Eleventh International Conference on Learning Representations*
45. Qi CR, Yi L, Su H, Guibas LJ. 2017. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems* 30
46. Scarselli F, Gori M, Tsoi AC, Hagenbuchner M, Monfardini G. 2008. The graph neural network model. *IEEE Trans. on Neural Networks* 20(1):61–80
47. Wu Y, Yan W, Kurutach T, Pinto L, Abbeel P. 2020. *Learning to manipulate deformable objects without demonstrations.* In *Robotics: Science and Systems XVI*, ed. AB Marc Toussaint, T Hermans, no. pap.065. N.p.: Robotics: Science and Systems Foundation
48. Longhini A, Welle MC, Erickson Z, Kragic D. 2024. Adafold: Adapting folding trajectories of cloths via feedback-loop manipulation. *Paper presented at the 4th Workshop on Representing and Manipulating Deformable Objects, IEEE International Conference on Robotics and Automation*
49. Cirio G, Lopez-Moreno J, Miraut D, Otaduy MA. 2014. Yarn-level simulation of woven cloth. *ACM Trans. Graph.* 33(6)
50. Sperl G, Narain R, Wojtan C. 2021. Mechanics-aware deformation of yarn pattern geometry. *ACM Trans. Graph.* 40(4)
51. Metaaphanon N, Bando Y, Chen BY, Nishita T. 2009. Simulation of tearing cloth with frayed edges. *Comput. Graph. Forum* 28(7):1837–1844
52. Coltraro F, Amorós J, Alberich-Carramiñana M, Torras C. 2024. A novel collision model for inextensible textiles and its experimental validation. *Appl. Mathem. Model.* 128:287–308
53. Blanco-Mulero D, Barbany O, Alcan G, Colomé A, Torras C, Kyrki V. 2024. Benchmarking the sim-to-real gap in cloth manipulation. *IEEE Robotics Autom. Lett.* 9(3):2981–88
54. Erickson Z, Clever HM, Turk G, Liu CK, Kemp CC. 2018. *Deep Haptic Model Predictive Control for Robot-Assisted Dressing.* In *IEEE International Conference on Robotics and Automation*, pp. 4437–44. Piscataway, NJ: IEEE
55. Antonova R, Yang J, Jatavallabhula KM, Bohg J. 2023. *Rethinking optimization with differ-*

entiable simulation from a global perspective. In *Proceedings of the 6th Conference on Robot Learning*, ed. K Liu, D Kulic, J Ichnowski, pp. 276–286, vol. 205 of *Proceedings of Machine Learning Research*, pp. 276–286. N.p.: PMLR

56. Ramos F, Possas R, Fox D. 2019. *BayesSim: Adaptive Domain Randomization Via Probabilistic Inference for Robotics Simulators*. In *Proceedings of Robotics: Science and Systems*, no. pap. 29. N.p.: Robotics: Science and Systems Foundation
57. Kemp CC, Edsinger A, Torres-Jara E. 2007. Challenges for robot manipulation in human environments [grand challenges of robotics]. *IEEE Robotics Autom. Mag.* 14(1):20–29
58. Isola P, Lim JJ, Adelson EH. 2015. *Discovering states and transformations in image collections*. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1383–91. Piscataway, NJ: IEEE
59. Bohg J, Hausman K, Sankaran B, Brock O, Kragic D, et al. 2017. Interactive perception: Leveraging action in perception and perception in action. *IEEE Trans. Robotics* 33(6):1273–91
60. Wang PC, Miller S, Fritz M, Darrell T, Abbeel P. 2011. *Perception for the manipulation of socks*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 4877–4884. Piscataway, NJ: IEEE
61. Shajini M, Ramanan A. 2021. An improved landmark-driven and spatial–channel attentive convolutional neural network for fashion clothes classification. *The Visual Computer* 37(6):1517–26
62. Liu Z, Luo P, Qiu S, Wang X, Tang X. 2016. *Deepfashion: Powering robust clothes recognition and retrieval with rich annotations*. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1096–104. Piscataway, NJ: IEEE
63. Duan L, Aragon-Camarasa G. 2022. A continuous robot vision approach for predicting shapes and visually perceived weights of garments. *IEEE Robotics Autom. Lett.* 7(3):7950–7
64. Willimon B, Birchfield S, Walker I. 2011. *Classification of clothing using interactive perception*. In *IEEE International Conference on Robotics and Automation*, pp. 1862–8. Piscataway, NJ: IEEE
65. Corona E, Alenya G, Gabas A, Torras C. 2018. Active garment recognition and target grasping point detection using deep learning. *Pattern Recognition* 74:629–41
66. Bouman KL, Xiao B, Battaglia P, Freeman WT. 2013. *Estimating the material properties of fabric from video*. In *Proceedings of the IEEE international conference on computer vision*, pp. 1984–91. Piscataway, NJ: IEEE
67. Runia TF, Gavrilyuk K, Snoek CG, Smeulders AW. 2020. *Cloth in the wind: A case study of physical measurement through simulation*. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10498–507. Piscataway, NJ: IEEE
68. Ha H, Song S. 2021. *FlingBot: the unreasonable effectiveness of dynamic manipulation for cloth unfolding*. In *Proceedings of the 5th Conference on Robot Learning*, ed. A Faust, D Hsu, G Neumann, pp. 24–33, vol. 164 of *Proceedings of Machine Learning Research*, pp. 24–33. N.p: PMLR
69. Colomé A, Pardo D, Alenya G, Torras C. 2013. *External force estimation during compliant robot manipulation*. In *IEEE international conference on robotics and automation*, pp. 3535–40. Piscataway, NJ: IEEE
70. Strese M, Brudermueller L, Kirsch J, Steinbach E. 2020. Haptic material analysis and classification inspired by human exploratory procedures. *IEEE Trans. Haptics* 13(2):404–424
71. Kampouris C, Mariolis I, Peleka G, Skartados E, Kargakos A, et al. 2016. *Multi-sensorial and explorative recognition of garments and their material properties in unconstrained environment*. In *IEEE International Conference on Robotics and Automation*, pp. 1656–63. Piscataway, NJ: IEEE
72. Yuan W, Dong S, Adelson EH. 2017. Gelsight: High-resolution robot tactile sensors for estimating geometry and force. *Sensors* 17(12):2762

73. Yuan W, Mo Y, Wang S, Adelson EH. 2018. *Active clothing material perception using tactile sensing and deep learning*. In *IEEE International Conference on Robotics and Automation*, pp. 4842–49. Piscataway, NJ: IEEE
74. Erickson Z, Xing E, Srirangam B, Chernova S, Kemp CC. 2020. *Multimodal Material Classification for Robots using Spectroscopy and High Resolution Texture Imaging*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 10452–59. Piscataway, NJ: IEEE
75. Kawabata S, Niwa M. 1989. Fabric performance in clothing and clothing manufacture. *Journal of the Textile Institute* 80(1):19–50
76. Wang H, O'Brien JF, Ramamoorthi R. 2011. Data-driven elastic models for cloth: modeling and measurement. *ACM Trans. Graph. (TOG)* 30(4)
77. Miguel E, Bradley D, Thomaszewski B, Bickel B, Matusik W, et al. 2012. *Data-driven estimation of cloth simulation models*. In *Computer Graphics Forum*, vol. 31, pp. 519–528. Wiley Online Library
78. Zheng D, Yao S, Xu W, Lu C. 2024. Differentiable cloth parameter identification and state estimation in manipulation. *IEEE Robotics Autom. Lett.* 9(3):2519–26
79. Sundaresan P, Antonova R, Bohgl J. 2022. *DiffCloud: Real-to-Sim from Point Clouds with Differentiable Simulation and Rendering of Deformable Objects*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 10828–35. Piscataway, NJ: IEEE
80. Duan L, Boyd L, Aragon-Camarasa G. 2022. Learning physics property parameters of fabrics and garments with a physics similarity neural network. *IEEE Access* 10:114725–34
81. Yang S, Liang J, Lin MC. 2017. *Learning-Based Cloth Material Recovery from Video*. In *2017 IEEE International Conference on Computer Vision*, pp. 4393–4403. Piscataway, NJ: IEEE
82. Chu V, McMahon I, Riano L, McDonald CG, He Q, et al. 2013. *Using robotic exploratory procedures to learn the meaning of haptic adjectives*. In *IEEE International Conference on Robotics and Automation*, pp. 3048–55. Piscataway, NJ: IEEE
83. Longhini A, Moletta M, Reichlin A, Welle MC, Held D, et al. 2023. *EDO-Net: Learning Elastic Properties of Deformable Objects from Graph Dynamics*. In *IEEE International Conference on Robotics and Automation*, pp. 3875–81. Piscataway, NJ: IEEE
84. Yan W, Vangipuram A, Abbeel P, Pinto L. 2021. *Learning Predictive Representations for Deformable Objects Using Contrastive Estimation*. In *Proceedings of the 4th Conference on Robot Learning*, ed. J Kober, F Ramos, C Tomlin, pp. 564–74, vol. 155 of *Proceedings of Machine Learning Research*, pp. 564–74. N.p.: PMLR
85. Lin X, Wang Y, Huang Z, Held D. 2022. *Learning Visible Connectivity Dynamics for Cloth Smoothing*. In *Proceedings of the 5th Conference on Robot Learning*, ed. A Faust, D Hsu, G Neumann, pp. 256–66, vol. 164, pp. 256–66. N.p.: PMLR
86. Huang Z, Lin X, Held D. 2022. *Mesh-based Dynamics with Occlusion Reasoning for Cloth Manipulation*. In *Robotics: Science and Systems XVIII*, ed. DS Kris Hauser, S Huang. N.p.: Robotics: Science and Systems Foundation
87. Tsurumine Y, Cui Y, Uchibe E, Matsubara T. 2019. Deep reinforcement learning with smooth policy update: Application to robotic cloth manipulation. *Robotics and Autonomous Systems* 112:72–83
88. Weng T, Bajracharya SM, Wang Y, Agrawal K, Held D. 2022. *FabricFlowNet: Bimanual Cloth Manipulation with a Flow-based Policy*. In *Proceedings of the 5th Conference on Robot Learning*, ed. A Faust, D Hsu, G Neumann, pp. 192–202, vol. 164 of *Proceedings of Machine Learning Research*, pp. 192–202. N.p.: PMLR
89. Hietala J, Blanco-Mulero D, Alcan G, Kyrki V. 2022. *Learning visual feedback control for dynamic cloth folding*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1455–62. Piscataway, NJ: IEEE
90. Avigal Y, Berscheid L, Asfour T, Kröger T, Goldberg K. 2022. *SpeedFolding: Learning Efficient Bimanual Folding of Garments*. In *IEEE/RSJ International Conference on Intelligent

Robots and Systems, pp. 1–8. Piscataway, NJ: IEEE

91. Lee R, Abou-Chakra J, Zhang F, Corke P. 2022. Learning fabric manipulation in the real world with human videos. *Paper presented at the 3rd Workshop on Representing and Manipulating Deformable Objects, IEEE International Conference on Robotics and Automation*
92. Petrík V, Kyrki V. 2019. *Feedback-based Fabric Strip Folding*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 773–778. Piscataway, NJ: IEEE
93. Ganapathi A, Sundaresan P, Thananjeyan B, Balakrishna A, Seita D, et al. 2021. *Learning Dense Visual Correspondences in Simulation to Smooth and Fold Real Fabrics*. In *IEEE International Conference on Robotics and Automation*, pp. 11515–22. Piscataway, NJ: IEEE
94. Xue H, Li Y, Xu W, Li H, Zheng D, Lu C. 2023. *UniFolding: Towards Sample-efficient, Scalable, and Generalizable Robotic Garment Folding*. In *Proceedings of The 7th Conference on Robot Learning*, ed. J Tan, M Toussaint, K Darvish, pp. 3321–3341, vol. 229 of *Proceedings of Machine Learning Research*, pp. 3321–3341. N.p.: PMLR
95. Seita D, Ganapathi A, Hoque R, Hwang M, Cen E, et al. 2020. *Deep imitation learning of sequential fabric smoothing from an algorithmic supervisor*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 9651–9658. Piscataway, NJ: IEEE
96. Blanco-Mulero D, Alcan G, Abu-Dakka FJ, Kyrki V. 2023. *QDP: Learning to Sequentially Optimise Quasi-Static and Dynamic Manipulation Primitives for Robotic Cloth Manipulation*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 984–991. Piscataway, NJ: IEEE
97. Tamei T, Matsubara T, Rai A, Shibata T. 2011. *Reinforcement learning of clothing assistance with a dual-arm robot*. In *2011 11th IEEE-RAS International Conference on Humanoid Robots*, pp. 733–738. Piscataway, NJ: IEEE
98. Zhang F, Demiris Y. 2022. Learning garment manipulation policies toward robot-assisted dressing. *Sci. Robotics* 7(65):eabm6010
99. Joshi RP, Koganti N, Shibata T. 2019. *A framework for robotic clothing assistance by imitation learning*. vol. 33, pp. 1156–1174. London, UK: Taylor & Francis
100. Seita D, Jamali N, Laskey M, Tanwani AK, Berenstein R, et al. 2022. *Deep Transfer Learning of Pick Points on Fabric for Robot Bed-Making*. In *Robotics Research: The 19th International Symposium ISRR*, ed. T Asfour, E Yoshida, J Park, H Christensen, O Khatib, pp. 275–290, pp. 275–290. Cham, Switz.: Springer International Publishing
101. Miller S, Van Den Berg J, Fritz M, Darrell T, Goldberg K, Abbeel P. 2012. A geometric approach to robotic laundry folding. *The Int. J. Robotics Res.* 31(2):249–267
102. Maitin-Shepard J, Cusumano-Towner M, Lei J, Abbeel P. 2010. *Cloth grasp point detection based on multiple-view geometric cues with application to robotic towel folding*. In *IEEE International Conference on Robotics and Automation*, pp. 2308–2315. Piscataway, NJ: IEEE
103. Luque A, Parent D, Colomé A, Ocampo-Martinez C, Torras C. 2024. Model predictive control for dynamic cloth manipulation: Parameter learning and experimental validation. *IEEE Trans. Control. Syst. Technol.* :1–17
104. Puthuveetil K, Wald S, Pusalkar A, Karnati P, Erickson Z. 2023. Robust body exposure (robe): A graph-based dynamics modeling approach to manipulating blankets over people. *IEEE Robotics Autom. Lett.* 8(10):6299–6306
105. Puthuveetil K, Kemp CC, Erickson Z. 2022. Bodies uncovered: Learning to manipulate real blankets around people via physics simulations. *IEEE Robotics Autom. Lett.* 7(2):1984–1991
106. Jia B, Pan Z, Hu Z, Pan J, Manocha D. 2019. Cloth manipulation using random-forest-based imitation learning. *IEEE Robotics Autom. Lett.* 4(2):2086–2093
107. Proesmans R, Verleysen A, Wyffels F. 2023. Unfoldir: Tactile robotic unfolding of cloth. *IEEE Robotics Autom. Lett.* 8(8):4426–4432
108. Erickson Z, Gangaram V, Kapusta A, Liu CK, Kemp CC. 2020. *Assistive Gym: A Physics Simulation Framework for Assistive Robotics*. In *IEEE International Conference on Robotics and Automation*, pp. 10169–76. Piscataway, NJ: IEEE

109. Garcia-Camacho I, Lippi M, Welle MC, Yin H, Antonova R, et al. 2020. Benchmarking bimanual cloth manipulation. *IEEE Robotics Autom. Lett.* 5(2):1111–1118
110. Clark AB, Cramphorn-Neal L, Rachowiecki M, Gregg-Smith A. 2023. *Household Clothing Set and Benchmarks for Characterising End-Effector Cloth Manipulation*. In *IEEE International Conference on Robotics and Automation*, pp. 9211–17. Piscataway, NJ: IEEE
111. Calli B, Singh A, Walsman A, Srinivasa S, Abbeel P, Dollar AM. 2015. *The YCB object and Model set: Towards common benchmarks for manipulation research*. In *2015 International Conference on Adv. Robotics (ICAR)*, pp. 510–17. Piscataway, NJ: IEEE
112. Garcia-Camacho I, Borràs J, Calli B, Norton A, Alenyà G. 2022. Household cloth object set: Fostering benchmarking in deformable object manipulation. *IEEE Robotics Autom. Lett.* 7(3):5866–73
113. Garcia-Camacho I, Borràs J, Calli B, Norton A, Alenyà G. 2022. Cloth manipulation and perception competition. *Paper presented at the 2nd Workshop on Representing and Manipulating Deformable Objects, IEEE International Conference on Robotics and Automation*
114. Pumacay W, Singh I, Duan J, Krishna R, Thomason J, Fox D. 2024. The colosseum: A benchmark for evaluating generalization for robotic manipulation
115. Zhou B, Zhou H, Liang T, Yu Q, Zhao S, et al. 2023. *ClothesNet: An Information-Rich 3D Garment Model Repository with Simulated Clothes Environment*. In *IEEE/CVF International Conference on Computer Vision*, pp. 20428–38. Piscataway, NJ: IEEE
116. Bertiche H, Madadi M, Escalera S. 2020. *CLOTH3D: clothed 3d humans*. In *European Conference on Computer Vision*, pp. 344–59. Springer
117. Bednarik J, Fua P, Salzmann M. 2018. *Learning to reconstruct texture-less deformable surfaces from a single view*. In *Int. Conf. on 3d vision (3DV)*, pp. 606–615
118. Gustavsson O, Ziegler T, Welle MC, Bütepage J, Varava A, Kragic D. 2022. Cloth manipulation based on category classification and landmark detection. *International Journal of Advanced Robotic Systems* 19(4)
119. Verleysen A, Biondina M, Wyffels F. 2020. Video dataset of human demonstrations of folding clothing for robotic folding. *The Int. Journal of Robotics Research* 39(9):1031–1036
120. Zhao J, Li J, Cheng Y, Sim T, Yan S, Feng J. 2018. *Understanding humans in crowded scenes: Deep nested adversarial learning and a new benchmark for multi-human parsing*. In *26th ACM Int. Conf. on Multimedia*, pp. 792–800
121. Mariolis I, Peleka G, Kargakos A, Malassiotis S. 2015. *Pose and category recognition of highly deformable objects using deep learning*. In *Int. Conf. on Adv. Robotics*, pp. 655–662
122. Thananjeyan B, Kerr J, Huang H, Gonzalez JE, Goldberg K. 2022. *All You Need is LUV: Unsupervised Collection of Labeled Images Using UV-Fluorescent Markings*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3241–3248. Piscataway, NJ: IEEE
123. Wagner L, Krejcová D, Smutnỳ V. 2013. Ctu color and depth image dataset of spread garments. *Center for Machine Perception, Czech Technical University, Tech. Rep. CTU-CMP-2013-25*
124. Schulman J, Lee A, Ho J, Abbeel P. 2013. *Tracking deformable objects with point clouds*. In *IEEE Int. Conf. on Robotics and Automation*, pp. 1130–1137
125. Chi C, Xu Z, Pan C, Cousineau E, Burchfiel B, et al. 2024. *Universal Manipulation Interface: In-The-Wild Robot Teaching Without In-The-Wild Robots*. In *Proceedings of Robotics: Science and Systems XX (to appear)*. N.p.: Robotics: Science and Systems Foundation
126. Yang PC, Sasaki K, Suzuki K, Kase K, Sugano S, Ogata T. 2016. Repeatable folding task by humanoid robot worker using deep learning. *IEEE Robotics Autom. Lett.* 2(2):397–403
127. Koganti N, Tamei T, Matsubara T, Shibata T. 2014. *Real-time estimation of Human-Cloth topological relationship using depth sensor for robotic clothing assistance*. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pp. 124–29. Piscataway, NJ: IEEE
128. Li Y, Hu X, Xu D, Yue Y, Grinspun E, Allen PK. 2016. *Multi-sensor surface analysis for*

robotic ironing. In *IEEE International Conference on Robotics and Automation*, pp. 5670–76. Piscataway, NJ: IEEE
129. Dai J, Taylor P, Sanguanpiyapan P, Lin H. 2004. Trajectory and orientation analysis of the ironing process for robotic automation. *International journal of clothing science and technology* 16(1/2):215–226
130. Estevez D, Victores JG, Fernandez-Fernandez R, Balaguer C. 2020. Enabling garment-agnostic laundry tasks for a robot household companion. *Robotics and Autonomous Systems* 123:103330
131. Estevez D, Fernandez-Fernandez R, Victores JG, Balaguer C. 2017. *Robotic ironing with a humanoid robot using human tools*. In *2017 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*, pp. 134–9. Piscataway, NJ: IEEE
132. Estevez D, Victores JG, Fernandez-Fernandez R, Balaguer C. 2017. *Robotic ironing with 3D perception and force/torque feedback in household environments*. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 6484–9. Piscataway, NJ: IEEE
133. Li S, Figueroa N, Shah A, Shah JA. 2021. *Provably Safe and Efficient Motion Planning with Uncertain Human Dynamics*. In *Proceedings of Robotics: Science and Systems XVII*, ed. MT Edited by Dylan A Shell, MA Hsieh. N.p.: Robotics: Science and Systems Foundation
134. Zhang F, Cully A, Demiris Y. 2019. Probabilistic real-time user posture tracking for personalized robot-assisted dressing. *IEEE Trans. Robotics* 35(4):873–888
135. Yamazaki K, Oya R, Nagahama K, Okada K, Inaba M. 2014. *Bottom dressing by a life-sized humanoid robot provided failure detection and recovery functions*. In *2014 IEEE/SICE International Symposium on System Integration*, pp. 564–570. Piscataway, NJ: IEEE
136. Canal G, Alenyà G, Torras C. 2019. Adapting robot task planning to user preferences: an assistive shoe dressing example. *Auton. Robots* 43(6):1343–56
137. Kapusta A, Erickson Z, Clever HM, Yu W, Liu CK, et al. 2019. Personalized collaborative plans for robot-assisted dressing via optimization and simulation. *Auton. Robots* 43:2183–2207
138. Erickson Z, Clever HM, Gangaram V, Turk G, Liu CK, Kemp CC. 2019. *Multidimensional Capacitive Sensing for Robot-Assisted Dressing and Bathing*. In *2019 IEEE 16th International Conference on Rehabilitation Robotics (ICORR)*, pp. 224–231. Piscataway, NJ: IEEE
139. Leidner D, Bartels G, Bejjani W, Albu-Schäffer A, Beetz M. 2019. Cognition-enabled robotic wiping: Representation, planning, execution, and interpretation. *Robotics and Autonomous Systems* 114:199–216
140. Dometios AC, Zhou Y, Papageorgiou XS, Tzafestas CS, Asfour T. 2018. Vision-based online adaptation of motion primitives to dynamic surfaces: application to an interactive robotic wiping task. *IEEE Robotics Autom. Lett.* 3(3):1410–17
141. Seita D, Jamali N, Laskey M, Tanwani AK, Berenstein R, et al. 2022. *Deep Transfer Learning of Pick Points on Fabric for Robot Bed-Making*. In *Robotics Research*, ed. T Asfour, E Yoshida, J Park, H Christensen, O Khatib, pp. 275–290, pp. 275–290. Cham, Switz.: Springer International Publishing
142. Li J, Sun W, Gu X, Guo J, Ota J, et al. 2022. A method for a compliant robot arm to perform a bandaging task on a swaying arm: A proposed approach. *IEEE Robotics Autom. Mag.* 30(1):50–61
143. Damayanti D, Wulandari LA, Bagaskoro A, Rianjanu A, Wu HS. 2021. Possibility routes for textile recycling technology. *Polymers* 13(21):3834
144. Fujii W, Suzuki K, Ando T, Tateishi A, Mori H, Ogata T. 2022. *Buttoning Task with a Dual-Arm Robot: An Exploratory Study on a Marker-based Algorithmic Method and Marker-less Machine Learning Methods*. In *2022 IEEE/SICE International Symposium on System Integration (SII)*, pp. 682–9. Piscataway, NJ: IEEE
145. Gries T, Lutz V. 2018. Application of robotics in garment manufacturing. In *Automation in garment manufacturing*, ed. R Nayak, R Padhye, pp. 179–197, The Textile Institute Book Series. Woodhead Publishing
146. Papoutsidakis M, Piromalis D, Priniotakis G. 2019. Advanced automation in textile indus-

try production lines. *International Journal of Engineering Applied Sciences and Technology* 4(5):504–507
147. Organization WH. 2015. *World report on ageing and health*. World Health Organization
148. Baek J. 2023. Smart predictive analytics care monitoring model based on multi sensor iot system: Management of diaper and attitude for the bedridden elderly. *Sensors International* 4:100213
149. Yu S, Lin K, Xiao A, Duan J, Soh H. 2024. *Octopi: Object Property Reasoning with Large Tactile-Language Models*. In *Robotics: Science and Systems XX (to appear)*. N.p.: Robotics: Science and Systems Foundation