Technical Report

IRI-TR-24-01

"Inside-out states of garments"

Pablo Jiménez Schlegl

February 2024

CSIC Consejo Superior de Investigaciones Científicas
Universitat Politècnica de Catalunya BarcelonaTech

Institut de Robòtica i Informàtica Industrial

Inside-out states of garments

Pablo Jiménez Schlegl

Institut de Robòtica i Informàtica Industrial (CSIC - UPC)
Llorens i Artigas 4-6, E-08028 Barcelona, Spain
e-mail: pjimenez@iri.upc.edu

Abstract

Cloth items such as garments can be potentially in infinite different configurations, due to their deformability. For the practical purposes of manipulation, however, it is possible to discretize the space of deformations into a set of equivalent states. This is pursued in the present work, where the complex problem of recovering the canonical configuration of partially (or completely) reversed garments is addressed. Without claiming to solve the hard perceptual and manipulative challenges of this type of task –which are also thoroughly described–, this work should rather be viewed as a pioneering effort to formalize the high-level strategies that aim at solving this canonical configuration-recovering task.

Keywords: deformable object manipulation, robotic garment manipulation, inside-out configurations, cloth state representation

---

1. Introduction

Domestic robots assisting in daily chores involving cloth handling will sooner than later be confronted with the problem of garment reversal, that is, cloth pieces partial- or totally inside-out, or partially outside-in[1] that have to be manipulated to regain their canonical rightside-out configuration. Be it for the presence in the household of kids, with their well-known entropy incrementation skills, or of people prone to undress in a hurry, or even users with some mobility constraints, or be it for the enigmatic shape twisting operations that take place in the spinning interior of the washing machine, the domestic robot may find that the garment it is picking up has parts of its inner side turned outside and vice-versa, as displayed in Figure 1. Proper arrangement of the garment means to turn the inside-out parts inside again, as well as the outside-in parts out. Sometimes also the inverse operation may be required, for example for washing or ironing delicate items, or for sun-drying dark dyed clothes (to avoid colors fading due to sunlight). However, despite being a common laundry-handling task, the scientific literature as for it being performed by a robot is extremely sparse. The reason is related to the overwhelming challenges such a kind of robotization has to face, as will be explained below.

The foremost challenge is the same that almost any kind of cloth manipulation is confronted with: the infinite number of potential configurations due to continuous deformations that a cloth item may attain. For the practical purposes of manipulation,

[1]completely outside-in coincides with completely inside-out.

Figure 1: Some examples in typical households of inside-out turned garments. Instances may be simple as a turned-out sock, or complex, involving intertwined garments.

however, it is possible to discretize the space of deformations into a set of equivalent states. A first coarse discretization is depicted in Figure 2: generic cloth states are shown which are described by every-day terms, common language concepts corresponding to qualitative experience. This figure does also display the actions that lead from one state to another.

The point in any state-space discretization for robot manipulation lies in accomplishing two objectives:

1. The state being identifiable and distinguishable from other states, and
2. The state having associated affordances, that is, providing the necessary hints for suggesting the subsequent action that the robot has to execute, as well as the required action parameters (e.g., visible grasp points).

We have to assume, for a meaningful instantiation of both objectives, that we are in a particular manipulation context, more specific than "do what has to be done with the next cloth item you detect". In other words, the robot is already focused on a finer granularity task, such as laundering, ironing, flattening, folding, etc., and perceives and acts accordingly. This assumption has led towards different types of discretization, which are briefly surveyed in Section 2.

What is exactly meant by inside-out (and outside-in) states, in the case of cloth? As, for simplification, cloth items are considered surfaces in space with negligible thickness, it is obvious that "inside" does not refer to the interior of the solid, as in the case of 3D objects, but to one of the two sides of these surfaces. In the case of simple rectangular or any other planar shape (e.g. circular or elliptical) there is either no meaningful distinction between the two sides (for example in towels or in scarfs), or this differentiation obeys to aesthetic reasons (printed patterns on sheets or embroidered tablecloth). In the latter case it certainly matters which side is facing up, while making the bed or dressing the table, but this is just an issue of the whole body orientation with respect to an external reference, the same as in the case of rigid objects. Thus, such simple surfaces will not be taken into account in this discussion, unless they form part of a more complex structure. Also non-orientable surfaces will remain out of consideration, although there actually

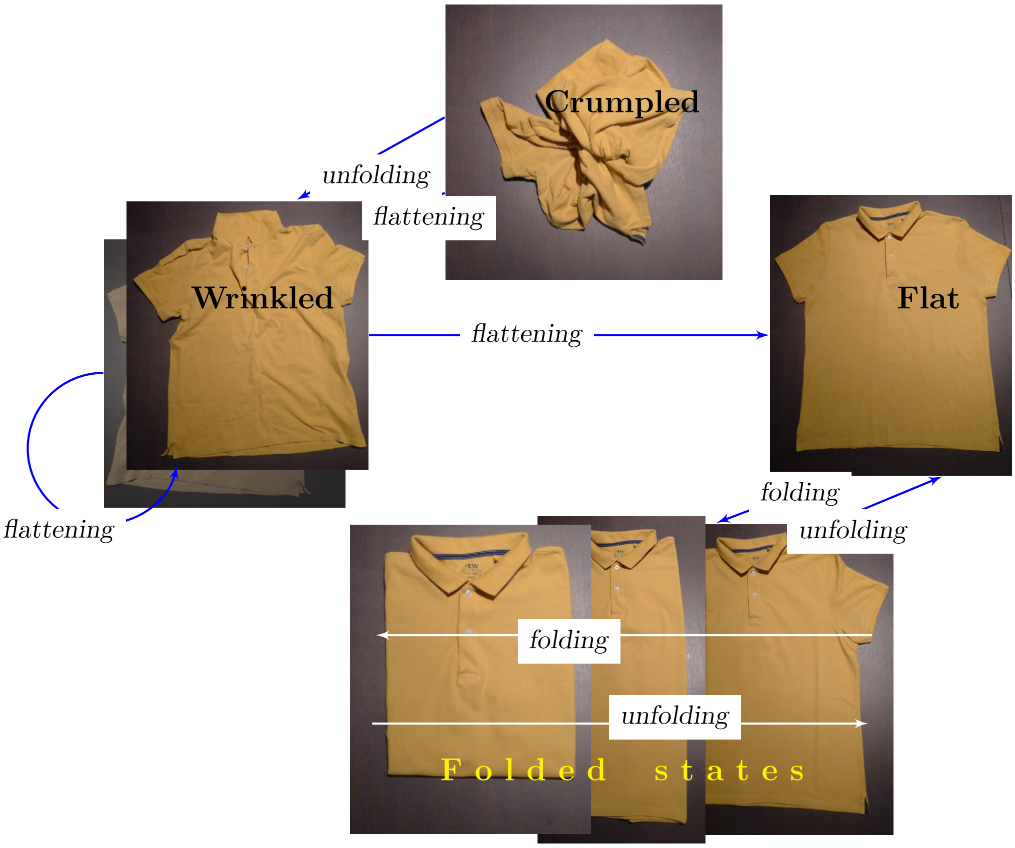

Figure 2: Actions that cause changes in the shape of cloth items. The circular arrow pointing from and to the Wrinkled state refers to flattening actions that gradually eliminate the wrinkles, without yielding the Flat state. The linear arrows across the Folded states refer to folding and unfolding actions between different discrete folded states. Unfolding in the air or from the Crumpled to the Wrinkled state may include auxiliary dynamic actions like shaking. Flattening from the Wrinkled to the Flat state may include the wiping, dragging and ironing variants, while if occuring from the Crumpled to the Wrinkled state only the dragging variant will be effective (in general it will refer to a first coarse flattening, which could also be just called extending*). For the sake of completeness, it should also be mentioned that the* crumpling *action leads from any state to the Crumpled state, but this action rarely arises in the domestic context (it is rather a side effect from putting cloth in a reduced space such as the inside of the washing maschine).*

exist Moebius loop scarfs and headbands. Instead, we refer to cloth parts that can be described as tubular or as hollow conic, and, as we are in fact talking about garments, it is straightforward to consider the inner side the one which is in contact with the human body, when the garment is worn correctly[2] by the user. Inside-out states are thus states where, due to deformation, the inside of these tubular or conical shapes has been partially or totally turned outside.

Inside-out states can obviously be represented by any geometrical model of the cloth item that differentiates between the inner and the exterior sides. A typical way to model cloth is via triangular or quadrilateral meshes (see [2] for a very recent contribution), and side distinction can be easily achieved by considering an "outward" normal for each one of the triangles (or, equivalently, a certain order –clockwise or counter-clockwise– of its vertices). This orientation has to be consistent across the model, meaning that adjacent triangles have the same outward side (i.e., if folded completely along the common edge, their outward (or inward) sides have to touch). Such a model has the advantage of being straightforwardly linked to graphical representations such as the ones used in simulations, but the assessment of inside-out states is computationally costly, as it requires to test the relative orientation of each triangle not only with respect to its neighbors but also distant patches of the model. We propose instead a simple representation where inside-out states can be rapidly determined and which also provides the action lines to restore the rightside-out states. This representation is described in Section 3 and the restoring or reversal actions in Section 4.

The problems encountered when facing the implementation of these actions in the real world are quite challenging. They are related to perception (identification of the state, detection of the grasp points for the reversal actions) as well as to grasping and manipulation. They are presented in Section 5, but their resolution falls out of the scope of this paper. Concluding remarks are shown in Section 6.

2. Related work

2.1. Cloth state discretization

We have already presented in the Introduction a first account of cloth states based on everyday experience and expressed in plain language. This discretization can be further refined, as has been done in different application fields. Previously to revise these more specific ad-hoc discretizations, the research presented in [16] has to be highlighted, which aims at a topological description of what they call macro-states, "set(s) of cloth configurations that can be manipulated in the same way, i.e., that have similar grasping affordances" (in their own words). The discretization of cloth states is obtained via a *stratification* of the topological space corresponding to the configurations in the plane of a set of significant points of the cloth. Each stratum is labelled with a binary vector which has as many elements as considered points, and the individual labels are computed as the signs of determinants of point coordinates, which express the relative position of a point wrt the (oriented) line defined by two other points. Their method is illustrated for

[2] *correctly* has now not only aesthetic connotations but also functional ones: try to button an inside-out worn shirt or to reach inside the pockets of inside-out worn trousers...

a cloth rectangle whose four corners are the considered significant points, and described also for the general n-point case. Figure 3 displays some configurations from different strata and they afford different cloth manipulation actions such as unfolding or turning over the whole rectangle.

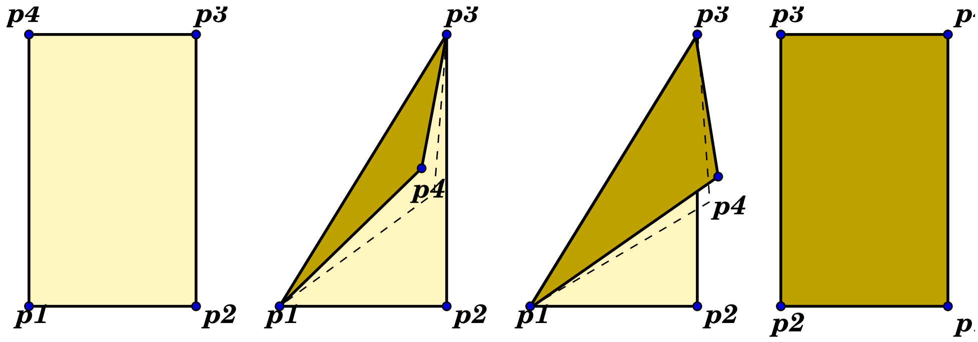

Figure 3: Configurations belonging to different strata, in the case of 4 points. The labels are given by the signs of the point coordinate determinants d_{123}*,* d_{124}*,* d_{134}*, and* d_{234}*, and are, from left to right: (+ + + +), (+ + - +), (+ + - -), and (- - - -). The back of the cloth is shown in a darker colour. The two central strata afford the same unfolding action by pinch grasping at* p4*, while the dashed lines suggest alternative configurations for the same strata.*

In the more specific context of garment folding, a set of canonical folded states can be defined via a finite set of fold lines and the resulting states of the (completed) individual folding actions. This is the case of the approach taken by Liu and Lin [12] who take the partition of the garment's surface induced by the user-defined fold creases as input to their folding sequencing algorithm. This is also the premise in Miller et al's [13] "g-folds" formalism, which is used in Figure 4 illustrating discrete fold states and the transitions induced by the folds. Although in our application there is no discretization of the state space by a previously established *location* of folds, we will see that the different states are defined by the *presence* of a special category of folds, namely inward or outward folds on tubular structures.

In flattening by sweeping and dragging applications (see, for example, [17]), states can be put in correspondence with the presence of wrinkles. Either just two states are defined (wrinkled/flat), or as many states of wrinkles to be removed. Another flattening application, namely ironing (see [3]), entails a discretization related to the partition of the garment into the regions laid out on the ironing table during the process (to which different ironing profiles apply).

Dressing assistance is another cloth handling application that admits considering different states of fulfillment. In [20] a state transition model describes the correct development of the dressing action as well as the dead end transitions towards two possible failure states, in the dressing of pants. Topological coordinates are employed in [10] to assess the progress of a T-shirt dressing operation on a mannequin: they allow to describe the correct placement of the T-shirt's neck wrt the mannequin's head and body, as well as of the sleeve on the mannequin's shoulder. In our work, we rely on simple geometric tests to determine reversed states of cloth on models, and on the presence of inward and outward folds, but considering a formalism based on topological coordinates could constitute an interesting alternative, worth further study.

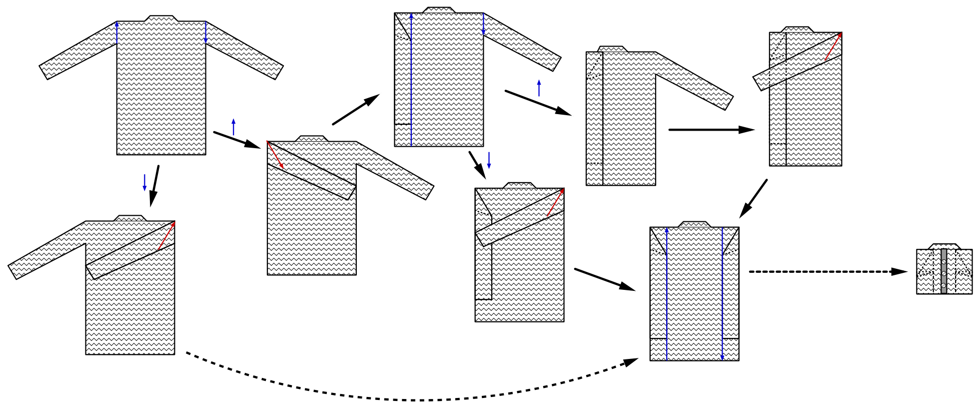

Figure 4: This graph represents some possible ways to fold a shirt (shown backface up). Symmetric as well as the final steps to obtain the terminal folded result are not displayed. We follow the convention described in [13]: the part to be folded is always on the left of the arrow. Blue arrows correspond to "g-folds" that affect the whole stacked geometry, whereas red ones correspond to folds that affect just the folded part in the last step. The point, as for our contribution, is to recognize the suitability of defining discrete states along a cloth manipulation.

The work that comes closer to ours, as they include inside-out states, can be found in [18]. Their texture- and feature-based algorithm aims at recognizing eight different states of socks, namely sideways, heel up, heel down, and bunched states, all in the rightside-out or inside-out variants. The choice of this restricted set of states is conditioned just for the context of laundry manipulation, with the goal of socks tidying and pairing. We aim at a more ambitious goal, as our domain extends over all type of garments and other items made of fabrics, and a higher number of states have to be considered. Nonetheless, their research points at a valuable direction as for the practical implementation, namely focusing on a feature-based state detection, which will be discussed in the second part of this paper.

2.2. Feature detection

As just said, feature detection will play a fundamental role for determining, via perception, reversed states of cloth items. It has been used for localizing specific grasping candidate points, as well as in cloth classification and state recognition [9, 8]. Besides generic features such as wrinkles or outline edges and corners, more specific features have been aimed, for grasping purposes, such as polo collars [14] (using a Bag of Features or Bag of Visual Words, BoVW, which combines appearance and 3D geometry features, the significant image areas are learnt via logistic regression and χ^2 SVM classification)[3], or shoulders of shirts and T-shirts or the corners of the waist of trousers and shorts in [5], for unfolding the garment in the air (using Hough trees).

As for classification, [19] follows a multi-level approach: low-level features, both global (Color Histogram, Histogram of Line Lengths, Table Point Feature Histogram, boundary) and local (Scale-Invariant Feature Transform (SIFT), and Fast Point Feature Histogram

[3]They later extended their method, for not specified application context, to detect not only collars, but also sleeves, hemlines or hips of different garment types (jeans, polo shirts, T-shirts, shirts and sweaters) [15]

(FPFH)) are used to determine and classify the mid-level *Characteristics* (i.e., cloth features like buttons, pockets, hemlines...). An SVM-based algorithm is employed to this end. Still in the mid-level, selection masks are determined from the database and used to assign the set of characteristics of the candidate cloth to the corresponding high-level class. For a particular case in which the category is already known (socks) but state recognition has to be performed [18] do also resort to a combination of 2D texture and shape-based features. The authors use two texture-based (MR8 filter bank and Local Binary Patterns (LBP)) and one shape-based feature (Histogram of Oriented Gradients (HOG)), and find out that the combination that performs best for inside-out vs. rightside-out classification is LBP plus HOG, while training the SVM classifier with a χ^2 kernel. As for state recognition, local detectors are trained with the appearance features to respond to image patches corresponding to the opening of the sock, the heel or the toe. Landmarks are placed on the centers of these patches and then compared with the parameterized models of each sock configuration. Finally, matching of pairs of socks is also accomplished by further using aditional cues related to size and color.

3. The representation

The human body can be roughly described geometrically as a polyarticulated tree of cylinders and some ellipsoids (the most obvious one of the latter being the head). This means that the repertoire of garment shapes needed to cover the body are restricted to simple surfaces, tubes, and hollow cones (or semi-ellipsoids)[4], which may also have holes, or combinations of them (including strings or bands to hold some parts together). Figure 5 displays some examples of such combinations for standard garments.

3.1. Reversed states of the basic shapes

The basic shapes that are subject to reversal actions are thus the tube and the hollow cone (or semi-ellipsoid), which represent either whole garments (a tube dress, a buff, or leg warmers for the tube, and a beanie or socks for the cone), or garments parts such as the trunk, the sleeves, the legs, etc. Thus, it is pertinent to describe the proposed representation first for these basic shapes.

The tube (whether perfectly cylindrical or frustoconical, this is not really relevant here) is represented by two oriented circles, each one corresponding to one end of the tube, and a line joining them, which we call seam line, without necessarily meaning that it corresponds to the actual seam of the corresponding garment (part). In the usual mesh-like representation of clothes for simulation, it is straightforward to label the set of edges corresponding to these circles and the seam line (the latter can be chosen arbitrarily, in fact we are only interested in its local behavior in the vicinty of each circle, as explained below). This extremely simple representation allows to describe geometrically the different states as for inside-out and outside-in reversals, as shown in Table 1. There we have considered those states involving the reversal of one or both of the tube extremes.

[4]Both shapes are considered equivalent for our purposes, we use them indistinctly, driven just for which is the most natural approximation. In the case of a pointed hood it is the cone, whereas in the case of the cup of a bra it is the semi-ellipsoid (unless you are considering Jean-Paul Gaultier's design for Madonna's bustier in the 1990 Blonde Ambition Tour).

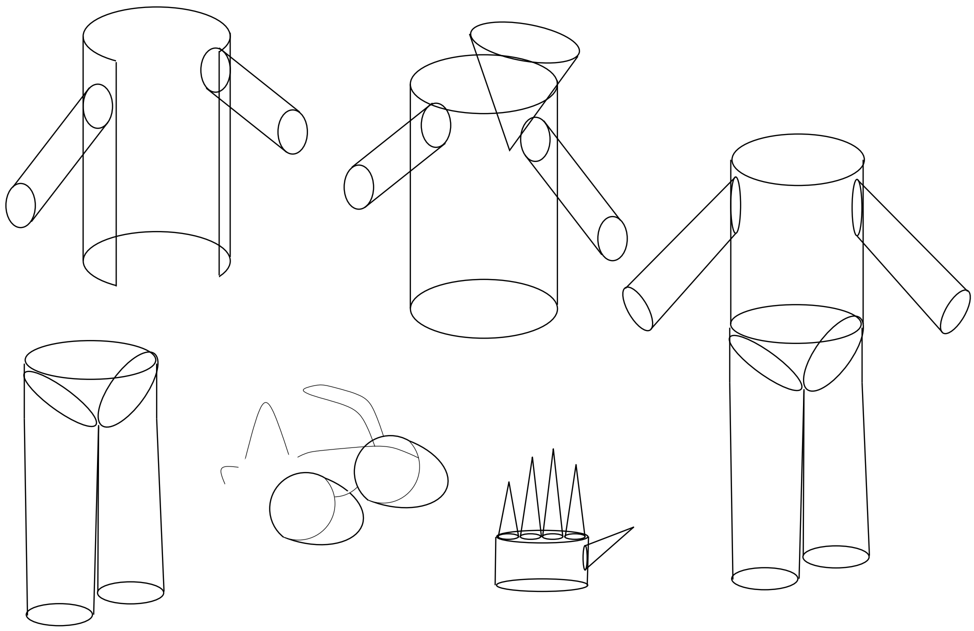

Figure 5: Garments as combinations of basic shapes. A shirt (in unbuttoned state, otherwise it would be assimilated to a pullover) consists in a combination of a simple surface (with holes) and two tubes, whereas a hooded pullover combines three tubes (one with holes) and a hollow cone, trousers two tubes and a cone with holes (aka slip), a bra two semi-ellipsoids (or cones) held together with strings, a glove a tube (or a cone) with holes (more than the typical two of the other shapes with holes) and five cones, and an overall as a cone with four holes (alternatively as a tube combined with a cone, both with two holes each, and of course in the zipped or buttoned state), and four tubes (sleeves and legs).

We have avoided the exhaustive depiction of symmetrical states, and also of those states resulting from repeated reversals (such as rolling up a sleeve). Observe that reversals may occur by flipping the inside of the tube outwards, or the outside inwards. Furthermore, the reversed end may be pulled along the tube's axis until the other end is reached. In the case of an outside inwards reversal, such pulling causes the inner side to become visible from the outside when the other end is reached. Following the criterion of defining states by what can be observed about them, rather than which course of actions have lead to them, we call inside-out states those where the interior of the tube has become visible from the outside, and the outside-in states those where exterior parts of the tube have become hidden.

Canonical or rightside-out		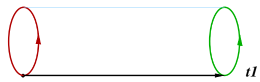	
	Outside-in	Inside-out	Both types
One tube end	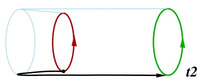	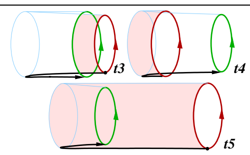	
Both tube ends	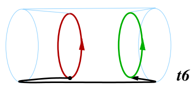	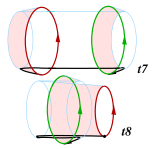	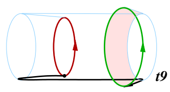
Completely reversed		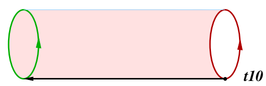	
	Inside-in*	Outside-out*	Both types
Both tube ends reversed item	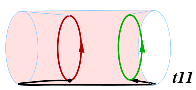	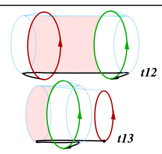	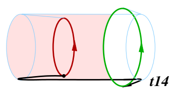

Table 1: Inside-out and outside-in reversal states of the basic tubular shape. The interior parts that have become visible from the outside due to reversal are shown shaded in pink. In all states but the first one, at least one of the tube extremes has been reversed, which is drawn as a red circle. If the tube has been previously reversed (partially or completely), then the states displayed in the lower row may be attained. These states have been marked with an asterisk, as the inside-in and outside-out state of the ends do not correspond to the canonical state. The tube ends are always the same size, they have just been drawn bigger or smaller to make the drawing more understandable.

The reversed states for the cone are displayed in Table 2.

In this ideal representation, a reversed end can be quickly discovered by simple geometric tests, as shown in Figure 6: the normal of the oriented circle is compared to the direction of the seam line, if their dot product is negative this end has been reversed. Executing this test on the openings of the tube and the cone (for the vertex of the cone, as there is no oriented circle, a different test is required, see Figure 6), plus some additional equally simple test (e.g., the circle-inside-the-tube test may be computed in $O(n)$

Canonical or rightside-out	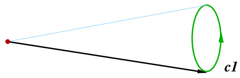		
	Outside-in	Inside-out	Both types
Vertex	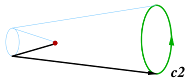	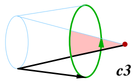	
Opening	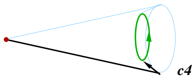	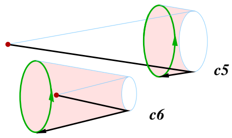	
Vertex and opening	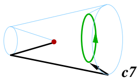	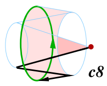	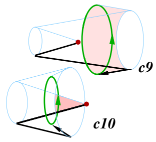
Completely reversed	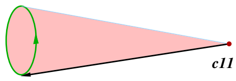		
	Inside-in*	Outside-out*	Both types
Opening	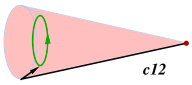	(Same state as $c3$)	
Vertex	(Same state as $c6$)	(Same state as $c5$)	
Vertex and opening	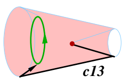	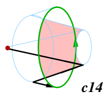	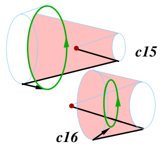

Table 2: Inside-out and outside-in reversal states of the basic conical shape. The interior parts that have become visible from the outside due to reversal are shown shaded in pink. Contrarily to the case of the tube, there is an intrinsic asymmetry in the case of the cone. Reversals affecting the tip of the cone can only be of the outside-in type. The cone opening is always the same size, it has been drawn smaller in the outside-in cases to make the drawing more understandable, actually it has to deform in order to fit inside the cone. Recall that this shape stands for actual cones, but also semi-ellipsoids or even cylinders closed at one end.

time, being n the complexity of the representation mesh, and this may be boiled down to constant by exploiting temporal coherence in successive simulation steps), all the different states may be discriminated.

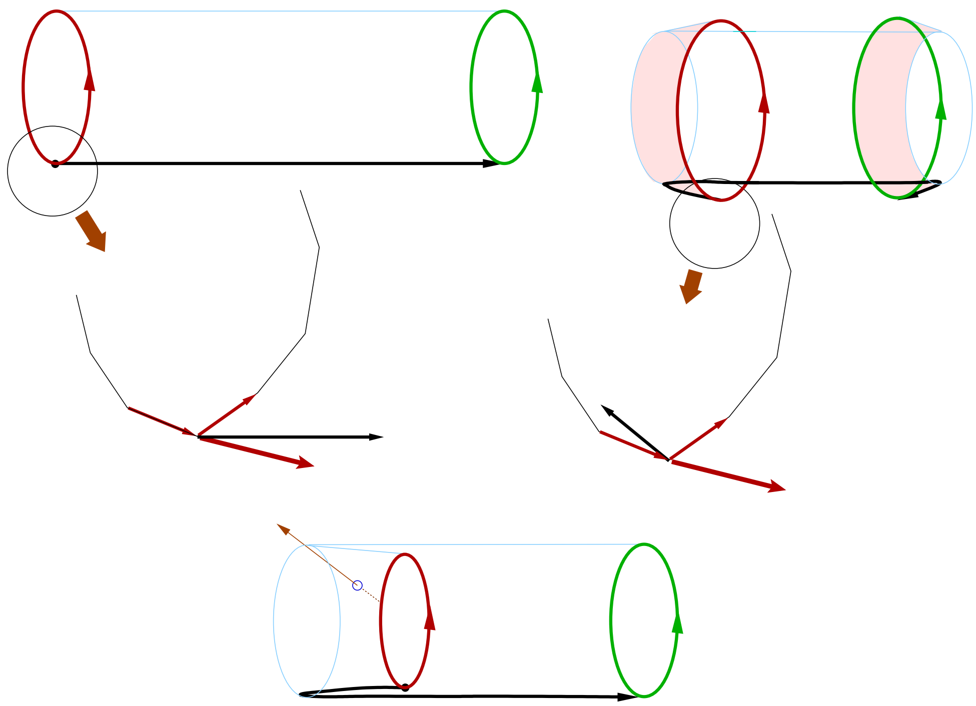

Figure 6: As soon as one of the ends is reversed, the seam line changes its orientation with respect to the normal of the oriented circle. The latter can be easily determined in a discrete representation of the garment by the cross product of the two incident edges, marked in red. As for discriminating the different reversed states, additional simple test can be used, such as determining possible intersections of an outwards directed ray from any vertex of the discrete representation of the oriented circle. In the case of the cone vertex, the normals of the adjacent faces are divergent, pointing outwards, in the canonical state, whereas they point inwards in the reversed state. This condition can be easily tested in the case where all of the incident edges to the vertex are convex (or concave, in the reversed state), via the intersection of the normal lines of two consecutive adjacent faces, which has to lie within the wedge defined by the semispaces the normals are pointing at, in the reversed state. The test is a little bit more involved when the model reproduces the existence of a crease

Figures 7 and 8 show all these states on real garments, a buff and a sock respectively. From these images (and from our everyday experience), it becomes evident that while such tests may work (possibly with some modifications) in simulations, the real stuff requires a different approach, mainly involving the detection of the inner side of garments, as well as discriminating the actual ends of the garment from those due to reversals (i.e., the circles drawn in light blue in Tables 1 and 2, we call them *reversal edges*). This question will be addressed (although not solved) in Section 5.1. Nonetheless, the representation fulfills the goal of discretizing the space of reversals, and can be consequently used for determining the necessary actions to reach a given desired state, as shown next.

It should be stressed that there is a type of outside-in deformations not contemplated in the previous discussion, as it does not involve one of the shape ends, but happens

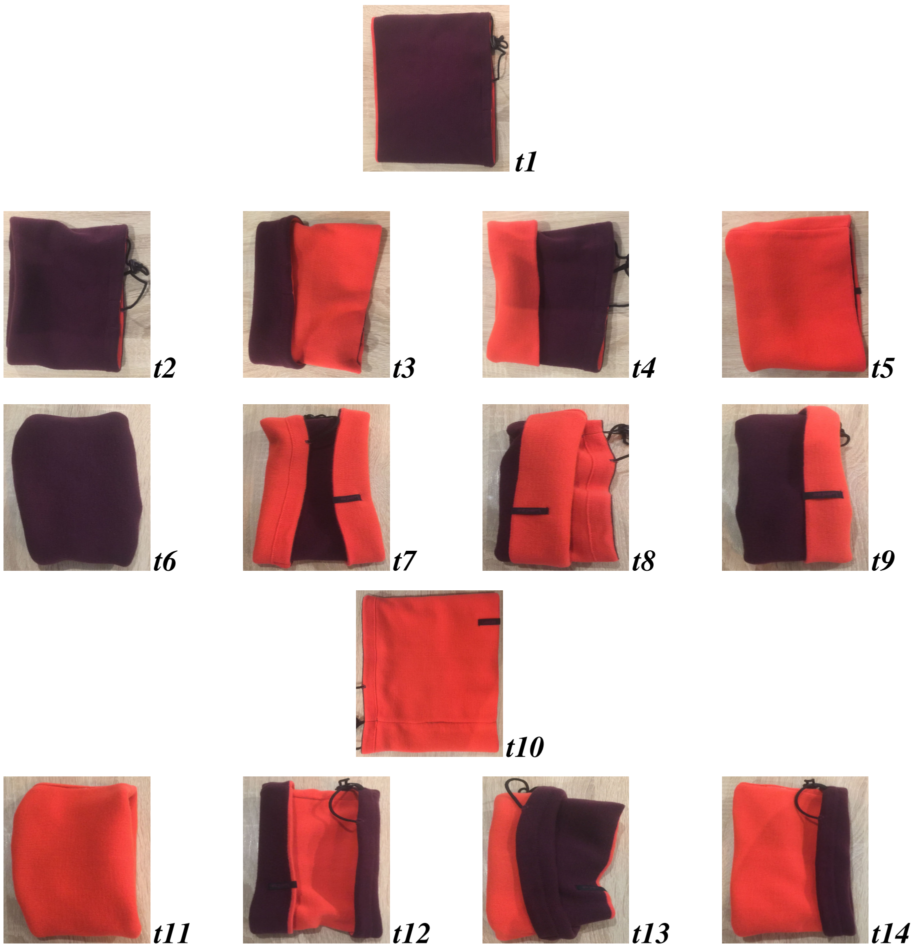

Figure 7: A buff in all the different states displayed in Table 1. As it is a reversible buff, the light red side has been arbitrarily chosen as the inner side. The canonical and the completely reversed states have been stressed with a green outline.

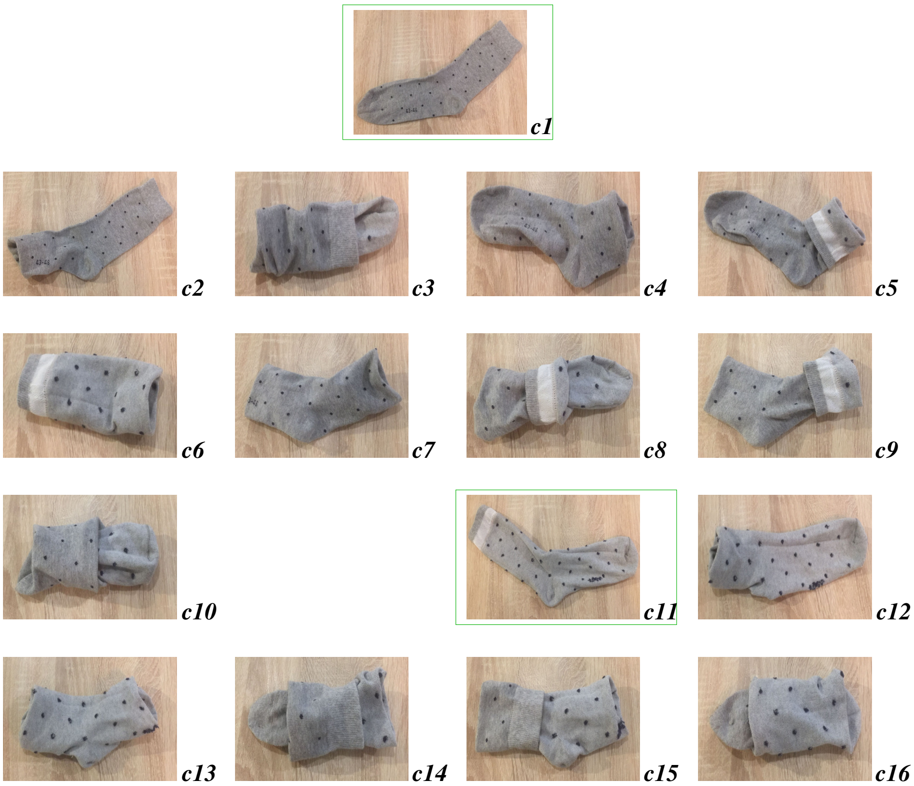

Figure 8: A sock reproducing the states shown in Table 2. The canonical and the completely reversed states have been stressed with a green outline.

anywhere between them. Figures 9 and 10 illustrate this type of deformations, which may also arise in the laundry manipulation context, possibly even more frequently than some of the states of Tables 1 and 2.

The geometric means to detect this situation in the context of modelling or simulation cannot resort anymore to the circles that represent the shape ends, nor the cone vertex, for obvious reasons. Nonetheless, the bending of the seamline (i.e., the direction changes along this line) provides a hint on where to localize the reversal edges, and perform similar tests on them as explained in Figure 6, which allow to distinguish these situations from simple shape bendings. Again, for real garment deformations (or realistic simulations), the assesment of these states involves rather the detection of reversal edges and the measurement of the overall length (such deformations involve an apparent shortening of the shape's length). See also the discussion in Section 5.1.1.

In the preceding discussion the basic shapes were assumed to have no additional holes. However, as was stressed in Figure 5, garments may also include basic shapes with holes, at which the connections of the "body" of the garment with the legs and/or sleeves occur. Also garments exist which are just such basic shapes with holes, see Figure 11.

Other examples of shapes with holes can be found in trousers, shirts or jackets with pockets. All the shapes with holes may present obviously the same set of inside-out (and outside-in) deformations or states of the basic tube and the basic cone, plus the deformations that result from portions of the shape traversing one or more of the suplementary holes, as shown in the same Figure 11. In composite shapes, common inside-out states in-

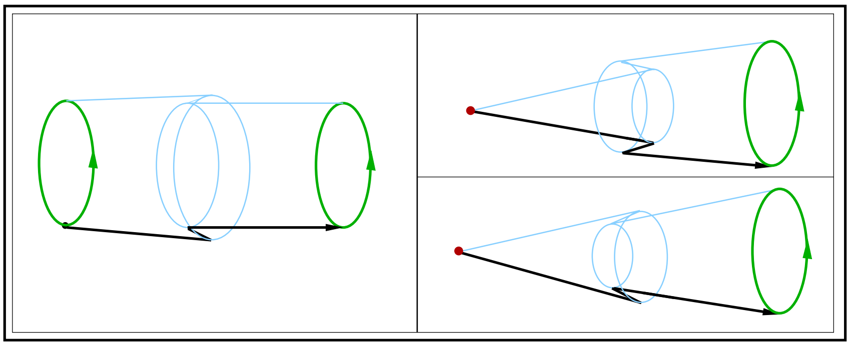

Figure 9: Outside-in deformations occurring not at the ends of the tube or hollow cone, but somewhere inbetween. In the case of the tube, just one state is depicted, as the assignment of the seamline direction and in circle orientations is arbitrary. However, it should be clear that two states have to be considered (as in the cone) as soon as the symmetry disappears, basically when the tube is a part of a composite shape.

Figure 10: Examples showing outside-in deformations occurring somewhere inbetween the shape's ends. Now, as the tube is a part of a composite shape (i.e., a leg of trousers), the symmetry disappears and two clearly different states have to be considered.

Figure 11: Basic surface, cone and tube with two suplementary holes, as well as a cone with four holes (which may also be regarded as combination of the two previous shapes). Below some real garments with such shapes are displayed, showing in some cases the additional inside-out deformation consisting in parts of the garment traversing these holes. One example of the last shape is the woman's one piece swimsuit.

clude other basic parts of the garment such as sleeves or legs traversing such suplementary holes, as illustrated in the next section.

3.2. Reversed states of composite shapes

Partial or total inside-out and outside-in states of shapes that combine simple surfaces, tubes and hollow cones, with or without suplementary holes, representing garments, such as the ones displayed in Figure 5 can also be represented as a discrete set of characteristic states. They can be seen as combinations of the states of basic shapes (Tables 1 and 2) plus protrusions of other parts of the garment through these supplemetary holes. Some of these combinations may not be impossible, but highly unlikely to appear in an unintentional or spontaneous way. For example, state $t2$ occurring on the shoulder end of a sleeve means to introduce part of the body of the jacket, shirt or pullover inside the sleeve, and this is not likely to occur in an everyday scenario of cloth manipulation.

Instead of exhaustively listing all the possible states for the different garment types (this is a laborious, although not impossible task, the variants of dressing up the human shape are limited, as already mentioned at the beginning of Section 3), we will concentrate on a specific garment, long trousers, to stress both the possible combinations of reversing the basic shapes and the limitations imposed by the fact of being attached to a common body, as well as the interactions between these elements. A subset of all the possible reversal states for this garment category is shown in Figure 12. The shown states are not an exhaustive depiction of the inside-out and outside-in states of trousers, for example the case where one leg is in state $t2$ and the other one at $t4$ is missing. Instead, we show some of the intermediate states between the canonical (trou1) and the completely reversed (trou13) states, as well as states that are dead-ends in this sense (trou5 and trou6), and states that result from the interaction of the basic shapes (trou3 and trou4). The recognition of anyone of these states allows to determine the appropriate sequence of actions required to obtain the final canonical or completely reversed state. If the goal state is trou13 and the current state is trou10, then the necessary actions that lead to trou11 and trou12 will be needed, and in the case that the recognized state is trou5, we will have to attain first the canonical state (trou1) before following one of the three sequences leading to trou13. In the next section, these reversal actions and their sequencing is analyzed more in detail.

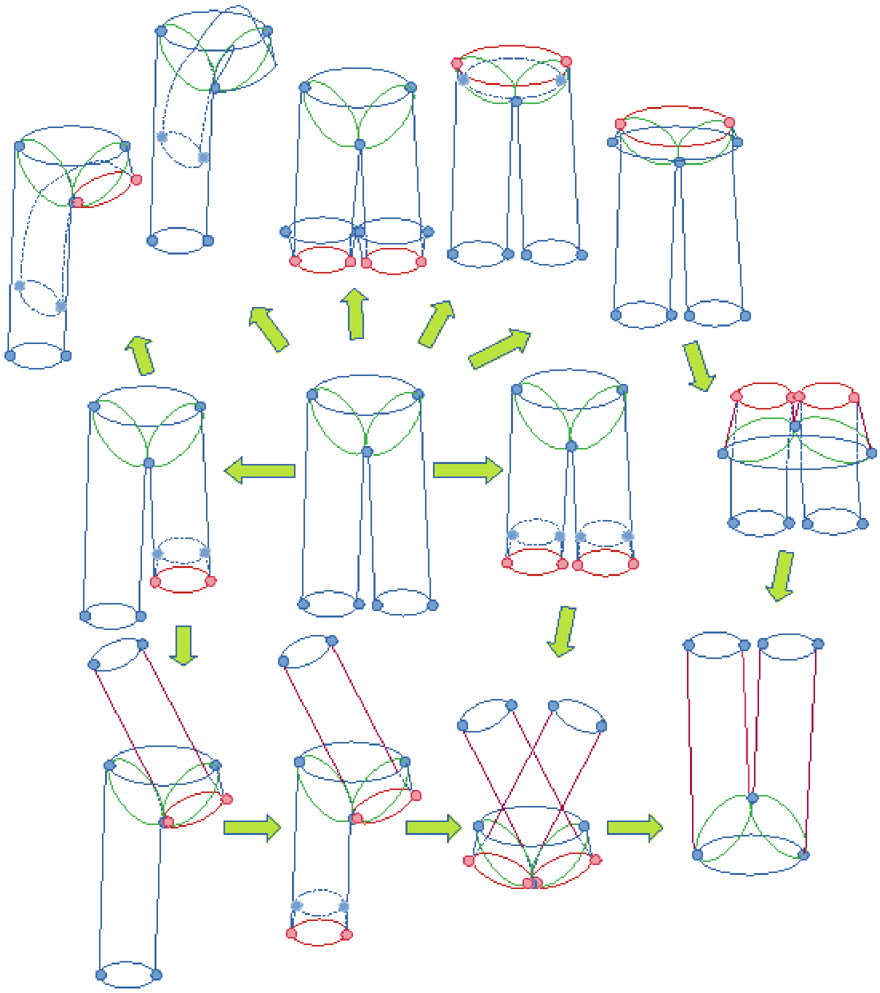

Figure 12: Possible states of trousers, as combinations of the states of the two basic tubes. The canonical state is trou1, and the completely reversed state trou13. Some states allow to attain both trou1 or trou13 indistinctly (e.g. trou9), others necessarily lead to either one or the other (e.g., trou13 cannot be reached from trou5 without attaining trou1 first). In states trou3 and trou4 one leg has been introduced into the other one, in this case a disentangling action (which in the case of trou3 is also a reversal action) is required. The particular notation of the states of the basic shapes in trou6 and trou7 just indicates the partial reversal of the corresponding end (as they share a common seam).

4. Reversal actions

4.1. Basic reversal actions

Without considering the real implementation details, which are dealt with in Section 5.2, three reversal actions may be distinguished:

- **end-flipping**, which is the local action of turning one of the ends of the basic shape inside-out or outside-in, or restoring a previously reversed end
- **pulling** the cloth along the axis, and
- **(supplementary) hole-traversing**, where parts of the same basic shape, or other parts of a composite shape are drawn through one of such holes .

Often, if not most of the time, the pulling action will implicitly conclude with an end-flipping. In composite shapes, pulling will also imply hole traversing in some cases (e.g., one leg of trousers which is completely reversed necessarily will have to traverse at some point the hole-union at the hip). Figure 13 provides graphical support to the definition of these elemental actions.

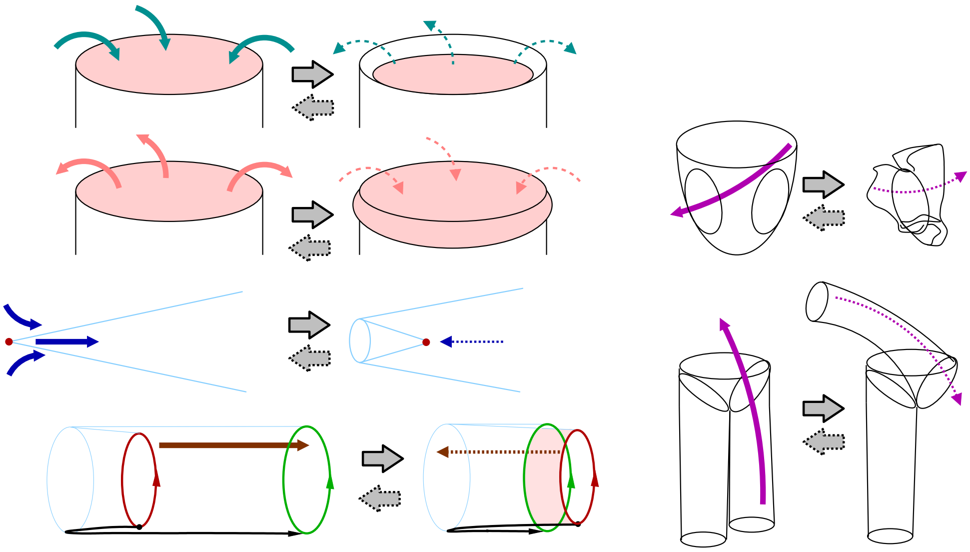

Figure 13: Elemental reversal actions. The end-flipping action either goes inwards or outwards (only inwards in the case of the vertex of the cone), and, if applied on a reversed end, may restore the non-reversed state (dashed arrows, we call these inverse *reversal actions). The pulling-along-the-axis action is parameterized by the direction, the magnitude of the displacement, and the application or pulling point (i.e., the end which is dragged). On the right, the hole-traversing actions, applied on a basic shape with holes and on a composite shape. Observe that the latter can be also viewed as the sequence of end-flipping and pulling-along-the-axis actions that, applied on the leg of the trousers, leads from its canonical to the completely reversed state (see Section 4.2).*

This is a somehow abstract definition of the actions, see Section 5.2 for practical considerations about their implementation.

4.2. Action sequencing

Elemental actions may be chained in order to reach any of the states in Tables 1 and 2 from any other one. Some state transitions are accomplished by a single action, others require to traverse a certain number of intermediate states. Figures 14 and 15 display the graph of possible state transitions for the basic shapes (without holes), any action sequence between two states can be found by determining a path between the corresponding nodes on this graph. Please observe that many more states are possible, such as those resulting from flipping not actual ends, but inward or outward folds; they have been left out to keep the graph reasonably simple.

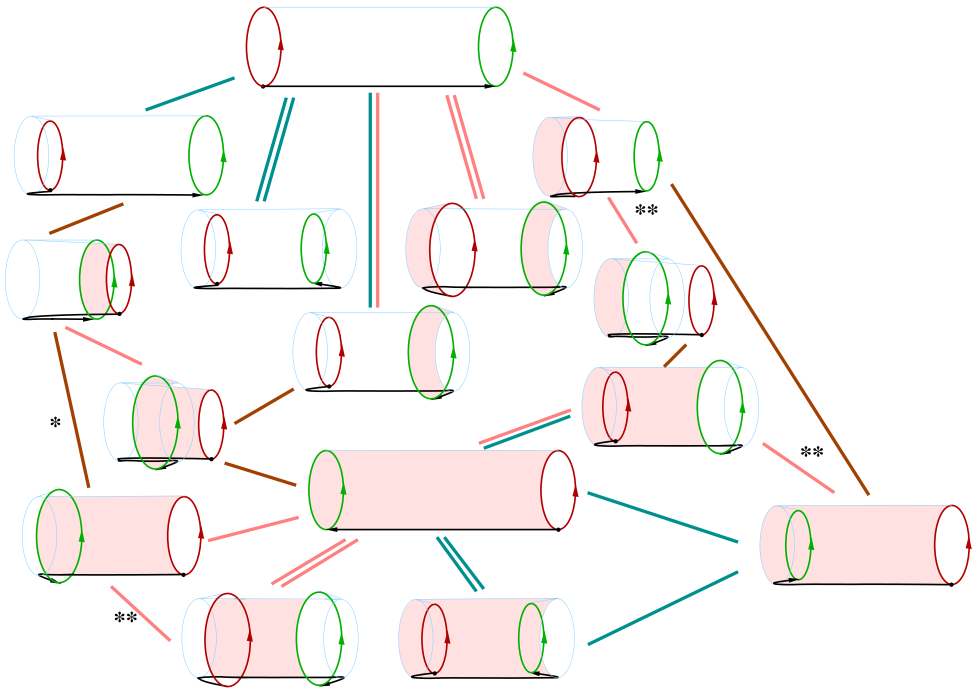

Figure 14: Graph of states (nodes) and elemental reversal actions (links) for the basic tube. The same code of colors as in Figure 13 has been employed here for the links denoting the corresponding elemental actions. Observe that it is not a directed graph, the same link stands for the elemental reversal action and its inverse. Some states are reached by a simultaneous application of reversal actions on both ends. (*) The two states linked by this arc are conceptually the same, but displaying them at different moments while pulling-along-the-axis allows to visualize better the following end-flipping action required for completion of the reversal.

Action sequences for composite shapes can be represented in the same way. In fact, Figure 12 already displays such a graph, in this case without distinguishing elemental actions, and assigning a direction that points towards the completely reversed state, as well as to some dead ends. The substitution of the arrows by the coloured arcs that correspond to the elemental actions is straightforward.

The state space and possible transitions is not large enough to justify the development of graph expansion procedures nor the application of sophisticated graph search techniques. There is no combinatorial explosion to be dealt with, just a few alternatives to get from one state to any other, which are already known beforehand. Moreover, the

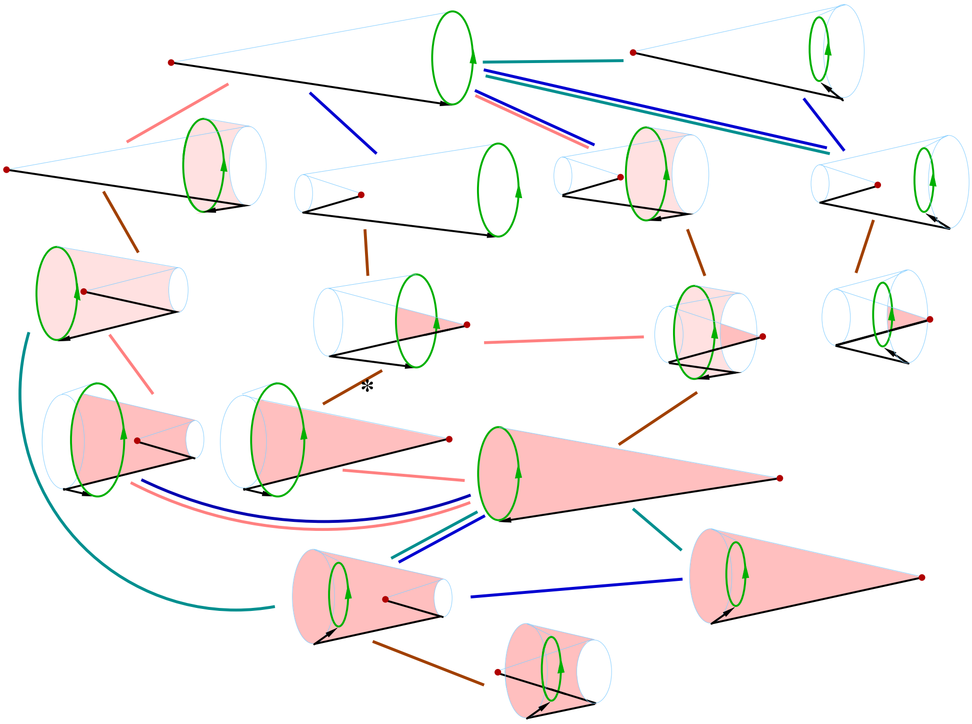

Figure 15: Graph of states (nodes) and elemental reversal actions (links) for the basic hollow cone. The same observations provided in the previous figure apply for this case as well.

manipulation goal will in most cases consist in attaining either the canonical or the completely reversed state. In the case of the tube, almost half of the possible states have a direct one-link connection to one of these two states (i.e., one or the other, the sets of states connected by a single link to these states are disjoint). The choice between the different courses of action can be grounded on the different degrees of difficulty associated to the elemental actions, which can be translated into arc (or link) costs. In general, the end-flipping action will involve a bimanual fine manipulation, which is more difficult than just pulling along the axis. However, depending on how the item is grasped, this end-flipping may automatically result from the pulling action, without requiring additional manipulation. These considerations (which are tackled in the next section) may imply a situational assignment of costs to the links, meaning that such costs depend on the specific manipulation situation the system is in, rather than fixed in advance. This is analogous, in the folding context, to the reachability constraints considered in [11], where the cost of performing a fold depends on the location of the robot (a mobile platform equipped with two arms) with respect to the garment, and the necessary displacement to reach the grasping point to initiate the folding.

5. Practical challenges in real implementation

We humans do perform cloth reversal operations practically without thinking. The general procedure involves picking up the cloth item, make a quick situation assessment (which may involve some sort of manipulation), introducing the hand through the body-side end of the reversed sleeve or leg, grasping its free-end (found using just touch), and pulling the reversed element back into canonical configuration while holding the main body of the garment with the other hand. Final shaking ensures elimination of residual deformations. Reversing items such as socks entails a similar operation, through (and holding) the only existing opening and grasping the toe instead. Smaller items such as baby clothing or gloves require dexterous manipulation. The difficulties encountered when trying to reproduce cloth-reversal skills with a robot are both perceptual and manipulation-related, as explained next.

5.1. *Perceptual problems*

5.1.1. *Situation assessment aka state recognition*

The visual recognition of any of the cloth reversal states described above has one basic requirement: to be able to distinguish the two sides of the cloth. This is far from being a trivial endeavour, as there is no single, always available feature that allows to make such distinction, no matter what garment is at hand. It is rather a set of features that have to be considered together to decide which is the inside, and even some manipulation, in the context of an active sensing scheme, may be necessary. Figure 16 identifies the most common features that can be used to this end, found on trousers in this case, but which can be extrapolated to other garment types.

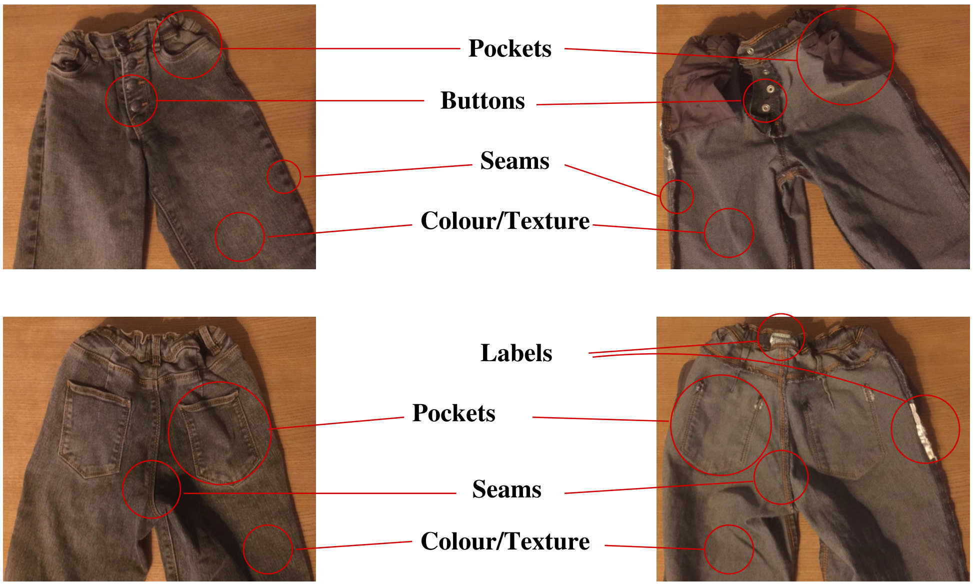

Figure 16: Front and back of right-side-out and inside-out trousers. The most common cloth-side identification features are shown.

These features and their idoneity are briefly discussed in what follows.

- **Surface color and texture.** In the textile and tayloring (or sewing) world, the two sides of a fabric are often referred to as the *right* and the *wrong* side, being the first the one which has to be outside (i.e., visible) in the finished garment or cloth item. Fabrics such as denim have a quite different colour on their two sides (due to the weave type, on one (the right) side the indigo dyed warp threads are predominant, whereas on the other the white weft is more visible). In such cases we are done:

this feature extends over the whole surface, and any reversed patch will have just a different colour. However, the system needs to know beforehand which colour corresponds to each cloth side. Similarly, printed woven or knitted fabrics display two quite different sides as well, being the right side the one with bright, definite colours, and the wrong side the one where the printed pattern appears faded. Some solid colour fabric types are woven or knitted in a way that quite distinguishable textures appear on each side, which can be appreciated even from standard vision alone, such as corduroy, velvet, or single knit fabrics. Other weave structures such as twill still display a slightly different texture on both sides, but the difference can only be appreciated via high resolution images. Some fabrics such as Cordura®have a finishing on the right side to make it waterproof and/or wear-resistant, which may in some cases make this side appear quite distinct from the wrong side, but in other cases the difference may only be told by touching. And finally there are fabrics with identical sides, as may be the case when the pattern is woven (not printed), as in plaid or gingham, or uniformily dyed fabrics as in the case of batik.

- **Seams.** In most cloth items, not only garments but also housecloth elements such as pillow cases, seams provide a powerfull feature for distinguishing the inside from the outside, as they are quite different on either side. As can be appreciated in Figure 16, their appearance on the outside is smooth, practically without relief, in this case with a single line of stitches (in other cases there may even be no visible stitches at all). They are quite different on the inside: they have a ridge-like shape, resulting from sewing together the two borders of the fabrics, and an intricated - albeit regular- stitch pattern, with longitudinal and transversal stitches. For a given garment type, they are usually located at the same places, which does also provide a hint for their localization if the garment is brought to a canonical position, such as the trousers in the mentioned figure: near or at the contour, and also along the crotch. However, in other configurations of the cloth item, wrinkles, foldings and self-occlusions in general may hide the seams. For this reason, the use of seams as a clue to determining the reversed state of a garment is advisable in the context of an active vision strategy, where multiple views of different parts of the garment can be gathered. It should also be noted that fashion trends may provide seams with quite a different appearance as expected, as can be seen in Figure 17.

- **Labels and tags.** In their most usual form, labels are small rectangular patches of fabrics sewn at the inner side of the garment. They are generally white (other colours are possible as well, in most cases different from the garment fabric's colour), and display some writing, the trademark of the garment (or a logo) and/or maybe the size or other information (e.g., washing and drying recommendations). In some cases the label is directly printed on the garment fabrics. Some years ago it was trendy to have the trademark label at the outside of the garment, and in the eclectic landscape of nowadays fashion there is still room for this option, as also shown in Figure 17.

- **Pockets.** There are two basic types of pockets, and both are illustrated in Figure 16: those made of a different, softer fabric, forming a small bag in the inside of the garment and accessed through a slit (front pockets in the figure), and those made of

Figure 17: This T-shirt on the left has rough, thick seams on the outside, due to ornamental reasons. However, the collar and the pocket clearly indicate that the garment is in a right-side out configuration. On the right, a label worn on the outside (Photo courtesy of LABELS, labelsfashion.com)

the same fabrics and sewn along their perimeter but the upper part on the outside of the garment (rear pockets). Pockets may also be shut with a flap (with or without buttons or checks) or a zipper. As different as the two types are, they are also quite different from the outside and the inside. And the two types have in common that the opening is accessible from the outside (the only exception being the inner pockets in jackets or coats, but in such cases there is no trace of the pocket on the other side of the garment, as they are completely hidden inside the inner lining).

- **Buttons, snap fasteners, hooks.** The knowledge about these fastening elements may also provide some hints about the state of the cloth item. Buttons have, besides their functionality, also ornamental purposes, which indicates where and how they are placed on the garment. They come in a huge variety of materials, shapes and colours: sewn flat buttons with two or four holes, shank buttons (the shank is a small loop on the back of the button, which allows it to be sewn), antlers and toggles, knot (or Chinese knot) buttons and frog buttons, metal buttons, snap fasteners, hooks and eyes... The garment closure requires both a button on one side and a buttonhole or a loop (or eye) on the other side of the closure. Sewn buttons appear on the outside of the garment, whereas their corresponding rear on the inside consists just in crossing stitches (which may be sometimes hard to detect due to the use of the same colour yarn than the garment). There are also shank buttons covered with the same textile than the garment, which makes them more difficult to detect using just colour, their detection is enhanced by the additional use of depth images. Metal buttons such as those in blue jeans (see also Figure 16) and most of snap fasteners are no sewn but have different parts one of which traverses the fabric, their appearance is quite different on either side of the cloth. Hooks and eyes such as those used on skirts, on bras, or to facilitate the closing of the zipper in trousers, are not meant to be seen and are located on the inside of the garment. It should also be noted that in garments such as shirts or blouses, both buttons and buttonholes appear along the front placket, which is usually made of the same fabrics as the garment, and may also present a different aspect form the outside and the inside.

- **Zippers and hook-and-loop (Velcro®) fasteners.** These clothing closures remain generally hidden when closed, either by the same garment fabrics (this applies also often to trousers buttons) or by a band of a different textile (generally in the inside). In open position, zippers appear partially covered by the garment fabrics in the outside, whereas in the inside (if not covered by a protective fabric tape) the zipper's tape is visible, as well as the stitches that attach it to the cloth item. Another hint is the position of the slider: the pull tab (or puller) is on the outside of the garment[5]. As for hook-and-loop fasteners, the common practice is to place (generally by sewing) the loop part on the wrong side (the inside) of the placket overlap and the hook part on the right side of the underlap. This is a valuable hint if the two parts can be distinguished: the hook part is generally lighter, smoother and more regular (if the image resolution is high enough, even the alignment pattern of the hooks may be appreciated).

- **Belt loops.** These are found on the exterior side of trousers and also of some jackets or coats.

- **Embroideries, bands, epaulettes, and other ornaments.** Garments and other cloth items may be decorated with different elements, embroidered, sewn or attached in other ways to their surface. For obvious reasons such ornaments will always be on the outside of the garment, whereas from the inside their location will either be concealed by the lining, or visible as a perimetral stitch pattern or loose end threads. Due to their ornamental purposes, they generally stand out from the rest of the garment, and should thus be easily detected by computer vision.

This list, despite its length, does not pretend to be exhaustive, but it should evince the richness and variety of the involved features. A computer vision system may certainly be trained to detect them, even under conditions of partial occlusion. However, this description has also stressed some of the difficulties associated to their consideration for inside-out state detection. Two main ideas should be extracted from this discussion:

- A single feature will in general be not enough to identify reversed parts on an arbitrary part of the garment's image, unless colour and or texture are sufficiently distinct on the two sides. Such cases include printed textiles (recall: faded colours on the inside), garments with lining (often with a shiny, silky appearance), and a few others. The more features are considered, the more robust may be the cloth side identification, although we have also seen pathological cases (such as exterior protuberant seams) where some features are rather detrimental for identification.

- High level knowledge about the usual arrangement of features in different garment types may be certainly useful in the localization and identification of the aforementioned features. Knowing in advance which garment is currently observed allows

[5]On reversible garments such as some types of jackets the puller may be repositionable on either side.

to exploit such knowledge. Conversely, classification of garments may rely on identifying these features. If the garment type is unknown, this suggests to enter an iterative classification - feature detection - state estimation process, starting from an initial hypothesis, and exploiting this knowledge, possibly within an active perception strategy. With time and experience (or sharing data and knowledge bases), such knowledge can be enriched and refined with more subtle facts, such as the existence of spare buttons sewn on an interior label in shirts, or of sewn trousers leg cuffs (don't try to unroll them!) to cite some examples.

As for the specific computer vision methods and learning algorithms suitable for detecting such features: due to different appearance of each type of feature, different approaches will be required for each category. In other words, a given garment image will have to be submitted to a set of specifically trained detectors. Methods such as the ones mentioned in Section 2.2 will have to be considered for their idoneity for every feature type. This will not be enough if the garment lies in an arbitrary, bunched configuration. In such a case, the system will have to engage an active sensing strategy, to reach at least an approximate spread out configuration, which allows to resort to high-level information about the expected relative location of the different features. Active sensing requires repeated regrasping, manipulation and perception of the garment, which is explained for the general (not inside-out) case in [8].

To proceed with the reversal of inside-out garments, once the garment state has been identified, the system has to identify at which points it has to be grasped for a successful reversal action. This is addressed in the two next sections.

5.1.2. Cloth holding grasp point and suitable opening detection

With the exception of some quite specific cases, the requisite for reversal actions will be the use of two arms: one of the hands holding the garment, while the other one pulls one extreme of the reversed garment part. Here we briefly examine the optimal (or a sufficiently good) location of the holding grasp. This location will depend on the cloth type and the reversed part of the garment, but a rule of thumb may be to grasp a right-side out patch close enough to the reversed part. For example, for a reversed sleeve, the shirt should be held at the corresponding shoulder, not at the other sleeve or the waist (this would certainly compromise the success of the reversal action). Figure 18 shows some suitable points where the garment may be held while reversing part of it.

Good candidates are located at the vicinty of the openings where the reversed part begins, e.g. the inward fold of a reversed sleeve. The location of such folded openings may also be necessary for another reason: they might constitute the entrance for the robot arm that has to perform the reversal. This will not always be the case, depending on the particular state and the objective of the reversal, the intrinsic openings (waist, collar, cuffs...) have to be chosen instead. In any case, distinguishing between the two types of openings is necessary, and may be accomplished by measuring (or comparing) the sharpness of the opening's edge: intrinsic openings are cloth borders and thus generally at most twice as thick as the fabrics, whereas folded openings are thicker (depending on the bending radius of the fabrics). In other words, the latter are smoother than the former, and this can be perceived via standard image processing (by comparing edge gradients, using similar techniques as for fold detection, which is the case in fact) or

Figure 18: Suitable holding points for reversal, marked in dark red. The green circles with a pink interior show possible locations of pulling points, which have to be reached from the inside of the leg or sleeve. (a) and (b) suitable holding and pulling points to turn a reversed leg rightside out again. (c) To complete the whole reversal of the pajama top, possible holding and pulling locations are shown. (d) A suitable point for turning the reversed sleeve rightside out. However, despite the opening where the arm performing the reversal has to be introduced is quite visible (e), the holding point is so far away from the pulling point, that the pulling hand becomes easily stuck before reaching destination. In such a case, a regrasping of a new holding point is required.

depth. In partially reversed garments, a third type of openings should be considered, namely intrinsic interior openings, such as for example the junction of the sleeve and the body. Such openings will in general be delimited by a seam. As for reversal actions, they may be considered as a particular case of folded openings, or taken into account for actions leading towards a completely reversed garment. This is illustrated in Figure 19.

Figure 19: Three types of openings that can be found on a partially reversed garment: intrinsic border openings (red arrows), intrinsic interior openings (pink arrow), and (temporal) folded openings (green arrows). For a reversed sleeve to be turned rightside-out, an intrinsic interior opening may be accessed from the outside, it behaves just as a particular case of a folded opening. Such an opening may also be accessed to reach a completely reversed configuration, as shown in Figure 18(c).)

5.1.3. Grasp point detection of the free end or tip

The robot arm that has to perform the reversal action faces a difficult perceptual challenge, namely to detect where to grasp, before it starts to pull. In most cases, the location of such a point remains hidden from the outside and has to be guessed. Taken a reversed sleeve as an example, and assuming the arm has already started to introduce itself through the opening, the cuff may be visible and the robot has to proceed just to the point where its own gripper tip becomes visible as well. In such case, it has to grasp the garment just at its current position (see the next section for the grasping challenges). Otherwise, the robot has to be equipped with touch sensors and rely on tactile sensing to detect when it has arrived to the sleeve's end (or, in the case of a sock, to the toe). This may not be trivial, as other parts of the garment surrounding the cuff may affect the sensor's readings. A tactile tracing of the interior of the sleeve along the motion of the arm may be furthermore affected by discontinuities that not necessarily indicate that the border of the cuff has been reached, as is the case with the shirt or blouse sleeve placket. If the length of the reversed part is known, this may be considered for estimating the current position of the gripper wrt the cuff (inextensibility can be assumed in most woven fabrics, but not for some knitted fabrics, where the uncertainty may be higher due to the deformations introduced by the arm itself). Once the free end of the sleeve (or the toe of a sock) has been reached, there is no priviledged grasp point, any point in the vicinty will do. Alternatively, instead of aiming at a unique grasp point, an iterative policy may be devised, alternating the grasping of arbitrary interior points along the sleeve, and short pulling actions, until the reversal is completed.

5.2. Grasping and motion in-contact

Grasping an exterior holding point corresponds to the standard cloth grasping problem, which counts already with a consolidated literature [1]. The most common mechanical pinch grasps aim at existing folds or create them by local deformation of the the fabrics by friction while closing the fingers in contact with the cloth. Clamp grasps may be suitable as well if the holding point is at the border of a garment's intrinsic opening. Other grasping principles have been used for manipulating fabrics, but rather in the manufacturing context (where also auxliar devices such as fixtures for holding the cloth may be used), and they cannot compete, as for the versatility required in domestic environments, with mechanical grasping. Mechanical grippers designed with the constraint of favouring cloth manipulation can be found in the literature (see [6, 4] for some recent contributions). For grasping the holding points, the only requirement is to provide a strong enough grasp to guarantee that the garment won't slip away while being manipulated by the other arm.

So far so good, but when it comes to grasp the point to be pulled to complete the reversal action, the system is confronted with the difficulties of grasping an interior point. And this point has to be reached, in the first place. This means for the robot arm to translate along and within a cloth tube, which may be quite narrow in some garments. The robot arm will be touching the cloth while it moves, and has to be able to reach its end (although maneouvres with regrasping of new holding points along the tube, while the other arm progresses along, can also be considered, as mentioned above), which imposes some constraints on the dimensions and external appearance of the arm and the gripper: the arm has to be long enough, both arm and gripper should have a smaller width than the cross-section of the tube (this is a hard constraint in the case of inextensible fabrics), they should avoid protuberances and rough surfaces that may cause the arm to get stuck inside the cloth, and the gripper, as the foremost extreme of the arm, should have an "aerodinamic" shape that facilitates avancing along a cloth tube which may present folds and twists (again avoiding sharp edges, rounded tips are preferred instead). This pointing shape of the gripper, which is suitable for sleeve or leg interior traversing, is not the most appropriate when the point has been reached and has to be grasped. Covering the inside of the gripper's fingers with a material with a higher friction coefficient would allow to reproduce the procedure we humans use to pinch interior points of a garment, together with a finger design that allows them to spread out. Socks and other conical shapes can be grasped at the tip in this way, whereas open-ended tubular shapes such as legs or sleeves require a previous reorienting of the gripper towards the wall of the cloth tube. Of course, other mechanisms may be devised, such as providing the pulling gripper's fingertips with retractable hooks or needles, but besides introducing the possibility of damaging parts of the garment, additional mechanisms and actuators complicate the gripper's design and make it more vulnerable to failure.

End-flipping, the reversal of a narrow strip adjacent to the end of a tube or the open end of a hollow cone may be just the result of a previous pulling action, but in many cases it will imply fine dexterous manipulation. Depending on the type of fabric, sometimes just an energic shaking will already undo this final fold. For certain types of fabrics such as most knitted cloth, in cases such as t_4 or c_5 (see Tables 1 and 2 respectively), the pulling gripper actually does not need to grasp any point, but just wipe over the fold to undo it. Similarly, inward folds (such as in t_2 or c_4) may be undone via an inside-out wiping of an open gripper. However, more persistent folds such as on the cuff of a shirt have to be undone by clamp grasping the folded end at two different locations (to adhere to Borràs' notation in [1], a 2PP grasp should suffice in most cases, but a 2LL one will be more effective) and performing an inwards or outwards flip, depending on whether we have an outward or inward fold, respectively. Ensuring such grasps is not a trivial endeavour, specially for inward folds.

Finally, it should be stressed that some reversal operations can be performed in a more efficient way, if human-like levels of dexterity are achieved. For example, a T-shirt can be completely reversed with a single movement by introducing both arms through the bottom hem, grasping each sleeve hem with one hand and by using both gravity and shaking, turn the piece inside out. Such type of manipulative shortcuts may be considered for robot arms whose dexterity and dynamic performance make them possible.

6. Conclusions

This paper does not pretend to solve cloth reversal by robotic manipulation. This is a very difficult problem, with a huge plethora of variants and particular cases, which has to confront quite involved perceptual problems. Instead, a conceptual framework is provided to formalize the different states of reversal of the basic (conical and tubular) and the composite shapes of garments and other cloth items. This framework not only introduces a necessary discretization that renders the world of cloth reversal configurations approachable, but also deals with the actions that lead from one state to the other. Furthermore, a systematic account of the practical perception and manipulation difficulties faced by cloth reversal has been described as well, paving the ground for further research which aims to solve these particular problems.

The reader, depending on their expertise, will be prone to stress one or the other aspect of the contribution; we, as the authors, consider necessary to highlight the following insights:

- Robot manipulation for cloth reversal is a very complex but not unsolvable problem. It is fundamental for assistive robotics to be effectively deployed in domestic environments.
- The state space of inside-out and outside-in reversals can be discretized, and actions leading from one state to the other can be formulated.
- Perception of reversed cloth parts is one of the main bottlenecks for an effective manipulation. Except the garments that present two uniformly distinct sides (colour, texture), cloth side identification must rely on considering multiple features described in Section 5.1.1.
- State assessment is further enhanced by detecting possible inward or outward folds. Gradient based edge detection techniques may allow to distinguish those temporal folded openings from the intrinsic ones.
- High-level knowledge about garments and expected localization of features for different types of clothing does not only help to classify correctly the observed item, but also to determine its state, including reversals.

- Active perception (lifting the garment, spreading it out to get additional or more informative views) will also be necessary to reduce uncertainty.
- Bimanual manipulation is unavoidable for cloth reversal actions. One robot arm holds the garment in place, grasped at a previously identified holding point, whereas the other one has to introduce itself inside the reversed part (or the part to be reversed), grasp the pulling point, and pull until completing the reversal action. Depening on the dimensions of the garment and the robot, such actions may entail regraspings of the holding point and relative displacements wrt the gripper that aims at the pulling point.

The challenges arising in perception are not unsurmountable, and in fact existing feature detection algorithms can be adapted to detect seams, pockets, labels, buttons, etc. Deep learning classification could theoretically be employed as well, but the lack of large databases on clothes turned (partial- or completely) inside out, and the absence of such images on the internet imply that such methods would only be applicable after a large dedicated experimentation for data gathering (gamification, collaboration and sharing could be resorted to, as well as devising some means of automated reversal - image aquisition procedures, or the use of simulation with hyperrealistic models). As in other robot cloth manipulation applications, benchmarking will be fundamental to measure progress and facilitate the sharing of research results [7].

Classification and state estimation go hand in hand (if the garment type is not known beforehand), and besides active perception, this process can be enhanced by high-level information on garment layout, as stated above, and this means building up an increasingly sophisticated knowledge base. Such a knowledge base would initially contain information about the basic geometry of the different garment types, further enriched with general facts such as pocket or seams location (e.g. jackets are generally designed for right-handed people, the interior pocket at chest height is thus on the left, but would appear on the right in a reversed state), and in latter stages be enriched with fine detail information such as that exterior pockets (lateral pockets on trousers or coat pockets) may have their openings covered by flaps, that (mainly military, but not only) shirts may have epaulettes on the soulders, or that buttons (or zippers) may also be found on the inner side of jackets or handbags, closing pockets.

The robot arms and grippers used for cloth manipulation and more specifically for performing reversal actions should certainly be fit for their purposes (i.e., their shape and dimensions not obstaculizing in-contact displacement in the inside of garments), without sacrificing their versatility for other laundry, or more general domestic tasks. A certain degree of dexterity will be required for grasping the interior pulling points, and more accurate performances for end-flipping or reversing parts of small cloth pieces.

New challenges will unfold when the specific perceptual and manipulative questions presented in this overview are effectively tackled and the corresponding solutions implemented. Every small advance, however, will bring the goal of a useful cloth management by assistive robots closer to fulfillment.

Acknowledgments

This work has been partially funded by the European Union Horizon 2020 Programme under grant agreement no. 741930 (CLOTHILDE), Project ROB-IN PLEC2021-007859 funded by MCIN/ AEI /10.13039/501100011033 and by the "European Union NextGenerationEU/PRTR" and Project CHLOE-Graph PID2020-118649RB-I00 funded by MCIN/ AEI /10.13039/501100011033

[1] Júlia Borràs, Guillem Alenyà, and Carme Torras. A grasping-centered analysis for cloth manipulation. *IEEE Transactions on Robotics*, 36(3):924–936, 2020.

[2] Franco Coltraro, Jaume Amorós, Maria Alberich-Carramiñana, and Carme Torras. An inextensible model for the robotic manipulation of textiles. *Applied Mathematical Modelling*, 101:832–858, 2022.

[3] J.S. Dai. Task analysis and motion generation for service robots: With reference to region segregation and path generation for robotic ironing. chapter 3, pages 30–50. 2012.

[4] Sònia Donaire, Júlia Borràs, Guillem Alenyà, and Carme Torras. A versatile gripper for cloth manipulation. *IEEE Robotics and Automation Letters*, 5(4):6520–6527, 2020.

[5] Andreas Doumanoglou, Andreas Kargakos, Tae-Kyun Kim, and Sotiris Malassiotis. Autonomous active recognition and unfolding of clothes using random decision forests and probabilistic planning. In *Robotics and Automation (ICRA), 2014 IEEE International Conference on*, pages 987–993, May 2014.

[6] Felix Von Drigalski, Daiki Yoshioka, Wataru Yamazaki, Sung-Gwi Cho, Marcus Gall, Pedro Miguell Uriguen Eljuri, Viktor Hoerig, Ming Ding, Jun Takamatsu, Tsukasa Ogasawara, and Jessica Beltrán. NAIST openhand M2S: A versatile two-finger gripper adapted for pulling and tucking textile. In *First IEEE International Conference on Robotic Computing, IRC 2017, Taichung, Taiwan, April 10-12, 2017*, pages 117–122, 2017.

[7] Irene Garcia-Camacho, Júlia Borràs, Berk Calli, Adam Norton, and Guillem Alenyà. Household cloth object set: Fostering benchmarking in deformable object manipulation. *IEEE Robotics and Automation Letters*, 7(3):5866–5873, 2022.

[8] Pablo Jiménez and Carme Torras. Perception of cloth in assistive robotic manipulation tasks. *Natural Computing: An International Journal*, 19(2):409–431, jun 2020.

[9] Pablo Jiménez Schlegl. Visual grasp point localization, classification and state recognition in robotic manipulation of cloth: An overview. *Robotics and Autonomous Systems*, 125:107–125, 2017.

[10] Nishanth Koganti, Tomoya Tamei, Takamitsu Matsubara, and Tomohiro Shibata. Estimation of human cloth topological relationship using depth sensor for robotic clothing assistance. In *Proceedings of Conference on Advances In Robotics*, AIR '13, pages 36:1–36:6, New York, NY, USA, 2013. ACM.

[11] Karthik Lakshmanan, Apoorva Sachdev, Ziang Xie, Dmitry Berenson, Ken Goldberg, and Pieter Abbeel. A constraint-aware motion planning algorithm for robotic folding of clothes. In Jaydev P. Desai, Gregory Dudek, Oussama Khatib, and Vijay Kumar, editors, *Experimental Robotics, STAR 88*, volume 88 of *Springer Tracts in Advanced Robotics*, pages 547–562. Springer, 2012.

[12] Honghai Liu and Hua Lin. Sequence trajectory generation for garment handling systems. *Applied Mathematical Modelling*, 32(6):1017 – 1026, 2008.

[13] Stephen Miller, Jur van den Berg, Mario Fritz, Trevor Darrell, Kenneth Y. Goldberg, and Pieter Abbeel. A geometric approach to robotic laundry folding. *International Journal of Robotic Research*, 31(2):249–267, 2012.

[14] Arnau Ramisa, Guillem Alenyà, Francesc Moreno-Noguer, and Carme Torras. Using depth and appearance features for informed robot grasping of highly wrinkled clothes. In *2012 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1703–1708, May 2012.

[15] Arnau Ramisa, Guillem Alenyà, Francesc Moreno-Noguer, and Carme Torras. Learning rgb-d descriptors of garment parts for informed robot grasping. *Engineering Applications of Artificial Intelligence*, 35:246 – 258, 2014.

[16] Fabio Strazzeri and Carme Torras. Topological representation of cloth state for robot manipulation. *Autonomous Robots*, 45(5):737–754, 2021.

[17] Li Sun, Gerardo Aragon-Camarasa, Simon Rogers, and J.Paul Siebert. Accurate garment surface analysis using an active stereo robot head with application to dual-arm flattening. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 185–192, May 2015.

[18] Ping Chuan Wang, Stephen Miller, Mario Fritz, Trevor Darrell, and Pieter Abbeel. Perception for the manipulation of socks. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2011, San Francisco, CA, USA, September 25-30, 2011*, pages 4877–4884, 2011.

[19] Bryan Willimon, Ian Walker, and Stan Birchfield. A new approach to clothing classification using mid-level layers. In *Robotics and Automation (ICRA), 2013 IEEE International Conference on*, pages 4271–4278, May 2013.

[20] Kimitoshi Yamazaki, Ryosuke Oya, Kotaro Nagahama, and Masayuki Inaba. A method of state recognition of dressing clothes based on dynamic state matching. In *System Integration (SII), 2013 IEEE/SICE International Symposium on*, pages 406–411, Dec 2013.