Chapter 5
Computational Aspects of Softness Perception

Massimiliano Di Luca[1], Marc O. Ernst[2]

[1] Centre for Computational Neuroscience and Cognitive Robotics, School of Psychology, University of Birmingham, Birmingham, UK

[2] Cognitive Neuroscience Department & Cognitive Interaction Technology-Excellence Cluster, University of Bielefeld, Bielefeld, Germany

1 Sensory Information About Softness

How do we choose a ripe avocado in a box of seemingly identical ones? How do we test a pillow in a shop to assess whether it will be comfortable once at home? How do we handle a container made of thin plastic, such as an open PET water bottle, that bends in as soon as we put our hand around it? The human brain is able to integrate the unique structural, motor, and sensory properties of our hands and body in order to effortlessly estimate the material properties of deformable objects.

Humans obtain information about material properties through several types of stereotypical manipulations, termed exploratory procedures (Lederman and Klatzky, 1996). During these interactions, sensory information about softness is obtained primarily from the tactile and proprioceptive sensory modalities (Tan et al. 1995; Srinivasan & Lamotte 1995). Signals from other sense modalities are also available and can contribute to softness perception. Consider, for example, the vibration produced by aggregate materials when perturbed (i.e., gravel), or the sound produced by the snow, or even the visual change in shape of a pillow when compressed. The vibrotactile, auditory and visual signals can provide information about compliance, at least to some extent. In this chapter, we will analyse the computational principles underlying the combination of multisensory signals during the interaction with deformable objects. Such an analysis facilitates the identification of the information processing mechanism that leads to softness perception.

Deformable objects can have a uniform material or be composed of several parts. In particular, the deformability of the object can differ between its surface and its interior volume. Springs, for example, are deformable objects that have rigid surfaces (see Chapter 1) so that there is no deformation of local shape at the contact point of the finger, but they are deformable as the external shape changes when force is applied. Such objects are not uncommon: we interact with such spring-like objects with rigid surfaces when we type on a keyboard. In a similar way, force feedback devices can render such spring-like objects, but such devices do not render tactile deformations that are typical of most everyday objects that we interact with. On the other hand, the surface of most compliant objects we encounter is deformable rather than rigid (i.e., pillows, beds, chairs, padded tools, driving wheels). This distinction is paramount as objects with deformable surfaces provide an additional source of information about their compliance, namely the local skin deformation of the finger (Srinivasan & LaMotte, 1995). There have been several attempts to create a haptic display that can render tactile information (Chapter 11), and some have even combined tactile and force-feedback displays (Scilingo et al., 2010).

Tactile information for objects with rigid surfaces specifies that the object is not compliant. This information is in conflict with force and position information that, when combined, leads to a different estimate of compliance (Fig. 5.1). The information available during the interaction with objects that have deformable surfaces can be made similar to the case of rigid surfaces, by either providing local anesthesia to the skin area in contact (Srinivasan & LaMotte, 1995), or by mediating the contact through a tool (e.g. a pair of pliers). Note that if the perceptual system has access only to force and position information, the brain has no choice but to estimate compliance by combining the force and position signals in the way shown by the formulas below.

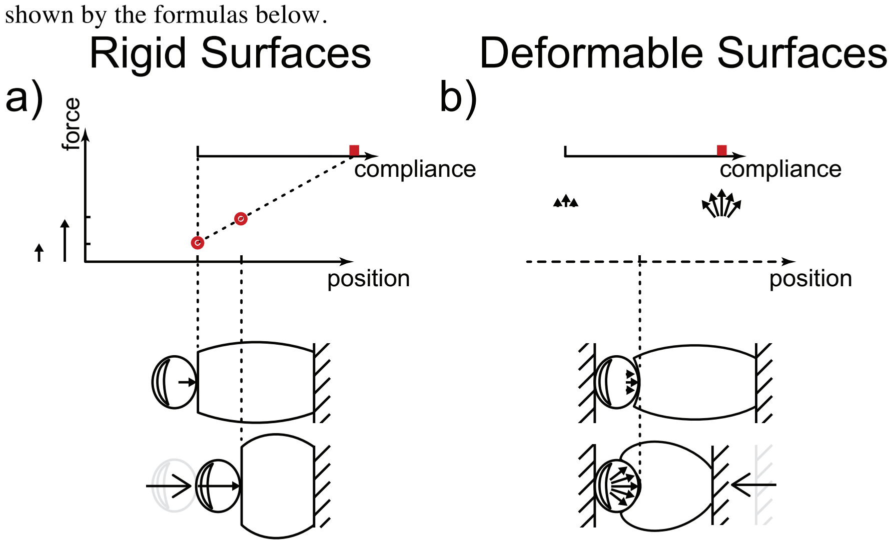

Fig. 5.1. Compliance information with objects having rigid and deformable surfaces. a) Compliance estimates obtained through the combination of position and force information according to Eqn. 5.1. The difference in position is divided by the difference in force (see Tan et al., 1995). b) The compliance of objects with deformable surfaces could be obtained from the pressure map, spread of the contact area and deformation pattern at the contact point, even when position information about indentation position is not available (Srinivasan & Lamotte, 1995; Bicchi et al., 2000). That is, tactile information about compliance is obtained in addition to the combination of force and position.

Object softness can be quantified for the two types of objects by identifying the amount of stress and the amount of object strain (see Chapter 1). Young's modulus (the ratio between force per unit area and strain) is a measure of stiffness (the inverse of compliance) of objects with soft surfaces. Young's modulus is a property of the material, whereas the stiffness changes with the width of the object. In the case of objects with rigid surfaces, the Young's modulus can be simplified to Hooke's law. Hooke's law states that the change in force (Δf) divided by the displacement Δp is the constant k defined as the object stiffness. The inverse of stiffness k is called compliance C, corresponding to position difference divided by force change:

$$C = 1 / k = \Delta p / \Delta f \qquad (5.1).$$

Compliance is the preferred term in this chapter as it is related to softness.

Several materials obey Hooke's law – i.e., compliance is constant throughout the interaction – as long as the force conditions do not exceed the material's elastic limits (Fig 5.2a). With non-Hookean materials, instead, force can change as a function of velocity and position, so it is necessary to update the relation between applied force and the change in indentation position in small increments during the exploration of the object (Fig 5.2b-d). As sensory information for softness perception is available at every instant during interaction, the brain needs to update the perceptual estimate of compliance from the sensory signals available at each time point. As we will discuss below, the time course of sensory signals is important in the perception of material properties as it carries information beyond compliance, e.g., about the rigidity of the object surfaces, about whether the object material is Hookean, or about whether the object is uniform.

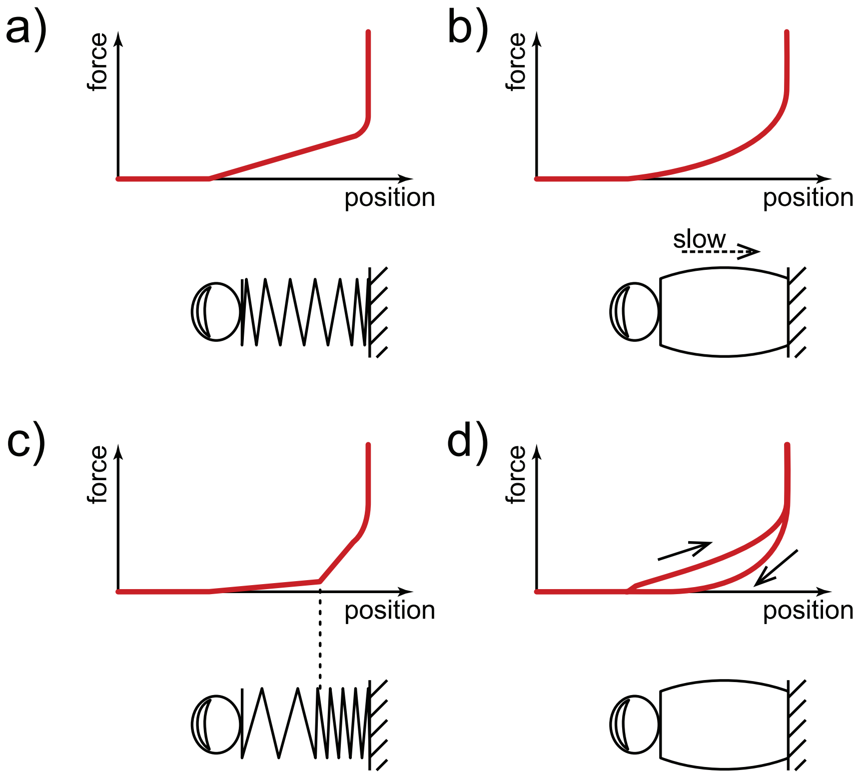

Fig. 5.2. Interaction with different types of objects that can lead to non-Hookean force-position relationships that become evident as deviations from linearity. a) Complete indentation of an object. b) Interaction with rubber or silicon specimen, where indentation velocity is low. c) Composite object (Di Luca, 2011) where the more compliant part is indented completely before the stiffer part. d) Material comprising a damping component where the force during loading and unloading movement does not match.

2 A Bayesian Model of Softness Perception

Different time-varying sensory signals specify the same value of compliance only within some error margin due to noise in the sensory system. That is, the multiple estimates do not necessarily agree and because of these precision limitations, the observer is faced with the task of inferring environmental properties from imperfect signals. The observer thereby makes a best guess to act upon and will then correct and update the estimate, as well as the action, as soon as more information becomes available. A solution to improve such a process is to combine information obtained from multiple sensory signals and from previous experience in order to predict the state of the world. It is particularly useful to use and combine signals from multiple sense modalities. For the manipulation of objects, useful sources of information can be derived not only from haptics (i.e., proprioception and active touch), but also from other sensory modalities such as passive touch, vision, and audition.

There is an increasing amount of empirical evidence supporting the idea that the brain processes sensory information in a way consistent with Bayesian Decision Theory (Knill & Richards, 1996). A Bayesian 'ideal observer' is a mathematical formulation of how perceptual estimates about a property of the world should be obtained and how these estimates should be used to optimally perform a decision (or action) given some cost function.

Bayesian Decision Theory offers a normative way to describe how observers should use information in order to form the most precise and accurate representation of the world, as all processes should be *statistically optimal*. In other words, the goal of the perceptual system should be to reduce uncertainty and, as such, this theory offers a benchmark against which human performance can be compared (Knill & Richards, 1996; Mamassian, Landy, Maloney 2002; Ernst, 2006). The key for this is the application of Bayes' theorem that uses probability distributions of a physical property in the environment, in this case object compliance.

To use statistically optimal computations like the Bayes' theorem, the brain should combine sensory signals that when taken alone do not carry information about the property of the environment in analysis. Furthermore, the brain should use all available sources of information to obtain a unique estimate of the physical property. Note that following Ernst and Bülthoff (2004), here we define the following terms:

- *Combination* indicates the processing of *complementary* sensory signals, where every component is necessary for an estimate of compliance. Combination of information leads to the likelihood component of the Bayes' formula (see Eqn. 5.2). Position and force information, for example, are both necessary in obtaining an estimate of compliance with an object having rigid surfaces (Fig. 5.1a). Combination of information does not lead to an increase in precision of the sensory estimate, as the noise of each of the two variables contributes to the noise of the final estimate.
- *Integration* indicates the use of two or more sensory signals that carry information about the same physical property. In Bayesian terms, each signal can specify a likelihood component of the Bayes' formula (see Eqn. 5.2). When interacting with objects that have deformable surfaces, information estimates of compliance can be obtained independently from haptic (using force and position, Fig. 5.1a) and tactile information (Fig. 5.1b). Assuming a relation between the global and the local deformation of the object, the two estimates are *redundant* and the brain can integrate them into one estimate (Fig. 5.3c). This generally will lead to an increase in the precision of the sensory estimate.

With these specifications, the Bayes' theorem states that the probability of a property "a posteriori" is proportional to the likelihood probability multiplied by the "a priori" probability, or:

$$p(property|evidence) \propto p(evidence|property) \times p(property) \quad (5.2).$$

When these values are expressed for every value of the environmental property, they can be represented as in Figure 5.3. Below, we analyze the components involved in this theorem when applied to compliance.

The first component we analyze is the function relating the probability of each possible sensory signal given a value of compliance, which is defined as the *likelihood function*. This is the probability that a state of the world could have generated the sensory signal that is available at the end of an interaction *p(sensory evidence | environmental property)*. In softness perception, the likelihood function is either obtained from information that directly specifies the compliance of the object (e.g., tactile information for objects with deformable surfaces), or by combining complementary signals (e.g., position and force information for objects with rigid surfaces).

According to the Bayesian theorem above, the likelihood function is combined with prior knowledge about the state of the world *p(environmental property)*, which is called the *a priori* probability distribution. The prior represents the statistics of the world with which we interact. For compliance, it represents the probability of encountering a compliance value even before any sensory information is available. In softness perception the prior could either represent the compliance experienced in prior interactions with a category of objects that can be identified visually (i.e. prior to touching it), or it could represent the frequency of compliances encountered during an individual's lifetime. We define the *conditional prior* to be the probability of a compliance based on the recognition of the object class, and we instead call the *statistical prior* what represents indistinguishably the compliances of all past objects encountered (see Fig. 5.3a-b). Note that the shape of the priors can greatly depart from a Gaussian distribution as it is determined by the statistics of the experienced environment, rather than the noise in the sensory processing, which is often Gaussian distributed.

The combination of likelihood function and prior distribution is proportional to the *posterior distribution, p(environmental property | sensory evidence)*. How does perception comes about? From the values of probability expressed at each level of the environmental property to be estimated (Fig. 5.3), the sensation of softness is usually thought to be determined from the maximum of this distribution. Because it is the Maximum of the *A Posteriori* distribution, this is also called the MAP estimate. A desirable result of combining the likelihood with the prior probability is that actions and perception will become more precise and also likely more accurate.

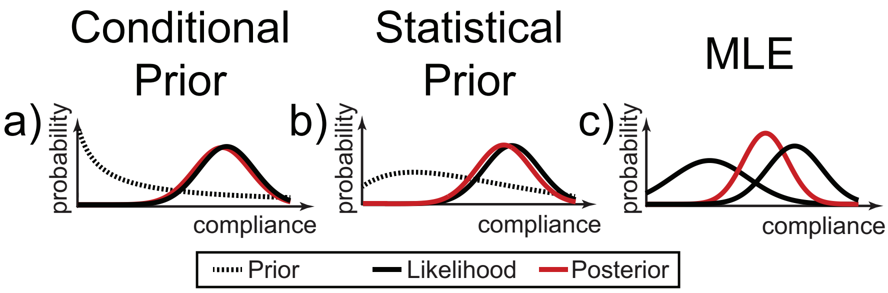

Fig. 5.3. Probabilistic representation of prior, likelihood, and posterior probabilities. a) Hypothetical *conditional prior* for the class of objects "avocado". Our experience of finding a hard avocado in a box exposed at the vegetable stand is higher than finding a soft one. Note that priors that convey little information (i.e., they tend to be flat) have less influence on perception. b) Hypothetical *statistical prior* that represents our overall experience with objects in our surroundings representing the fact that is unlikely that we could come in contact with objects with very high or very low compliance, but very low compliance (i.e. hard objects made of metal or wood) is relatively more probable. c) Maximum Likelihood Estimation with two Gaussian-shaped likelihoods deriving from independent redundant estimates of compliance where prior knowledge is non-influential (i.e. flat, not represented).

In many cases the prior probability of encountering some property of the world is equally distributed, and thus has little or no influence on the perceptual estimation process. Often, however, we encounter more than one source of sensory information about a particular world property, providing input to the estimation process. Under the assumption of a common cause, likelihood functions obtained from these different sensory signals should be combined into one final estimate of the probability of the state of the world. The application of the Bayesian framework has been very successful in describing how humans integrate redundant sensory information to obtain a more precise percept of the environment. When *a priori* knowledge can be assumed not to exert an influence on perception (i.e., when the range of compliance tested is small and each stimulus is presented an equal amount of times), the prior probability is "flat" and thus it is possible to consider the final estimate as being obtained through the integration of only the likelihood functions.

When the percept is assumed to be determined from the maximum of the combined likelihood function, we call it the Maximum Likelihood Estimate (MLE). In many instances it has been shown that human performance to integrate multisensory signals is close to the statistical optimum according to MLE (Ernst & Banks, 2002). With the assumption that the noises of such functions are all Gaussian in shape and independent, the integrated MLE estimate is a weighted average of the individual uni-sensory estimates, with weights proportional to the inverse variance of the unisensory distributions (the inverse of the variance is the precision and in the cue integration literature is also often termed reliability, Backus & Banks, 1999).

The result of MLE is that the integrated estimate is more reliable than either of the two components (Fig. 5.3c). Such a scheme has been applied to the perception of compliance with visual and haptic information about the position of the fingers (Kuschel et al., 2010; Di Luca et al., 2011) and for multiple contact points with the object (Di Luca, 2011; Plaisier et al., 2012). The perceptual consequences of integration are evident based on the magnitude of the perceptual estimate when information in the different modalities contains a small conflict, i.e., when visual and haptic information indicates a different amount of indentation (Kuschel et al., 2011) and when fingers are in contact with an object which has different compliances at the two contact points (Di Luca, 2011).

An important outcome of the integration of redundant sources of sensory information is that the uncertainty of the perceptual estimate will be reduced. In order to prove that for compliance perception there is indeed such a reduction in uncertainty, researchers have been comparing performance in discriminating material properties with one and two sources of sensory signals available. If performance with two sources of information is higher than the best single-source, the brain must be taking advantage of the redundancy in the sensory estimates (the signature of sensory integration, Ernst, 2006). For compliance perception, the outcome of multimodal integration has been shown to be close to the statistical optimum (maximal reduction in sensory noise) in some cases (Di Luca, 2011), but not in others (Kuschel et al., 2010; Cellini et al., 2013). Deviation from optimality seems to occur when conflicts are introduced by the experimental manipulation.

One conflict between redundant sensory signals that can lead to an overweight of information, and can also be ascribed to the detection of discrepancy, is with objects composed of different materials (Bergmann-Tiest & Kappers, 2009). Information from the deformable surface (sensed through touch) and the overall indentation of the object (sensed by combining force and position) can be made unrelated. With such objects, there is a discrepancy in the compliance estimate sensed from two sensory channels.

A second type of conflict that can lead to non-integration, which does not follow the MLE scheme, is when the visual information about the amount of indentation is manipulated (Kuschel et al., 2010; Cellini et al., 2013). This situation is discussed in detail in Chapter 2. Results suggest that integration uses a fixed set of weights and optimality is limited to the “natural” interaction. Such an outcome could be explained by hypothesising that the conflict is detected, thus increasing the chance that signals are not actually coming from the same external event.

In both cases of conflict (tactile-haptic and visual-haptic), information processing would have to balance the costs of loosing access to the individual estimates against the benefits of improving perceptual precision. Such balance depends on whether the conflict is present at the level of the final sensory estimate of compliance (tactile-haptic) or at the level of the individual complementary force and position components (visual-haptic). In the following section we will evaluate what these two possibilities entail.

3 Redundancy in Softness Cues

Integrating sensory information is only beneficial if the estimates are related to the same environmental aspect. Wrongly integrating sensory information leads to inherently erroneous estimates due to conflation. For this reason, integration should occur automatically only once the perceptual system has established the correspondence between estimates (Roskies, 1999; Ernst & Bülthoff, 2004). It has been shown that temporal and spatial coincidence (Thurlow & Jack, 1973), as well as structural similarity between the signals (Parise et al., 2012), is used as an indicator for the perceptual system to consider multimodal information as being related to the same source. Spatial and temporal offset between multisensory signals can, in fact, prevent integration (Witkin et al., 1952; Bresciani et al., 2005). In more complex situations, however, such as during manipulation of deformable objects, spatial and temporal coincidence can be overridden by an inference-like process (Duda, Hart & Stork, 2001) that determines whether sensory signals carry information about the same physical property (Helbig & Ernst, 2007).

In the perception of softness there are several ways in which sensory information could be redundant. With the notable exception of touch signal about objects with deformable surfaces (Srinivasan & LaMotte, 1995; Bicchi et al., 2000; Fig. 5.1), most sensory signals carry information only related to the indentation or to the force – they do not specify compliance directly. For example, information about the amount of finger movement is carried redundantly by visual and kinesthetic information, but auditory and tactile signals (vibrations and spread of the contact area) also provide some information (contributions of sound is discussed in Chapter 4, vibration in Chapter 3, area spread in Chapter 11). Information about resistive force is carried by touch and proprioception modalities and again other modalities could provide information about force (i.e., vision see Chapter 2 and Cellini et al., 2013). Thus, integration could occur at the level of position and force information, even before a compliance estimate is obtained. The brain could obtain a perceptual estimate of compliance in two ways (see Kuschel et al., 2010):

- Integration of force and/or position estimates (integration before combination). The brain proceeds by first integrating all the position signals into a unique position estimate, and all force signals into one force estimate. Only in a second step are position and force combined to obtain a compliance estimate (Fig. 2.2a, Fig. 5.4a&c).
- Integration of compliance estimates (combination before integration). The brain obtains separate estimates of compliance for pairs of signals providing force and position information. In a second step, the various compliance estimates are integrated into a coherent one (Fig. 2.2b, Fig. 5.4b,d).

It is still not entirely clear which strategy the brain adopts, and if so, whether it adopts the same strategy under all circumstances. It is possible that the selection depends on which integration strategy is more advantageous in terms of the reliability of the final estimate (see equations below). In an attempt to answer this question we will focus on two examples of interaction with rigid surfaces: First, we will consider the case of proprioceptive and tactile information, where force information is available in each modality (Fig. 5.4a,b), and then we will consider the case of visual and haptic information, where position information is available in both modalities (Fig. 5.4c,d — see Chapter 2 for an in-depth analysis of the visual-haptic case). In the analyses we will assume that the non-indented position of the object and an unloaded null force are used as references for compliance estimates. In such cases, position p and force f could be used in the formulas instead of Δp and Δf. This might not be the case as additional information is available, for example when the distance between the object and the hand is manipulated or when object is made to vibrate orthogonally to the contact point. If no other information is available about the location of the first contact, because of the vibration the boundary of the object should be sensed outside the object (i.e., earlier when approaching the object). The sensed indentation difference between null force and peak force should thus appear to be larger. Following this logic, vibration should have the effect of making object appear softer if the boundary is sensed before the interaction takes place (see Chapter 3 for a discussion about vibrotactile information and Chapter 9 for more information about boundary crossing).

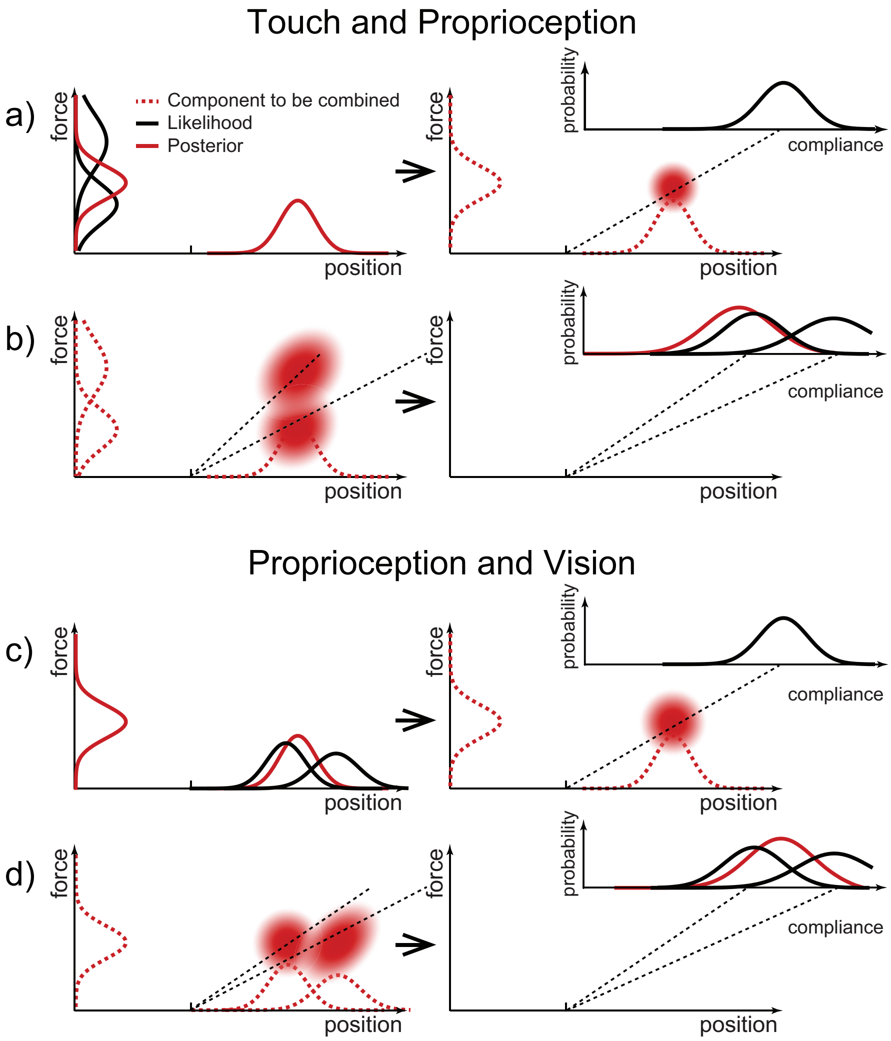

Fig. 5.4. Probabilistic representation of how integration and combination processes could be ordered to obtain a coherent estimate of compliance of objects with rigid surfaces from multiple sensory signals. a) Integration of force information provided by tactile and proprioceptive modalities (integration before combination). The integration is followed by combination of the posterior estimate of force to obtain a compliance estimate. b) Combination of force signals with proprioceptive position into separate estimates of compliance (combination before integration). In a second step, the two estimates of compliance are integrated into a coherent one. c) Integration of position information provided by vision and proprioceptive modalities (integration before combination; here tactile information about force is not represented). The next computation step is the combination of the force and position to obtain a compliance estimate. d) Combination of force with position information, provided by proprioception and vision, to obtain two compliance estimates which are subsequently integrated (combination before integration).

3.1 Touch and Proprioception

When force information about the indentation of an object with rigid surfaces is conveyed by both proprioceptive and tactile modalities (f_p and f_t), and vision is precluded, the reliability of force (and consequently of compliance), sensed through proprioception, depends on the mechanical configuration and properties of the arm and hand. Chapter 9 discusses how the reliance on position or force information depends on the distance of the contact point between the participant and the object. Similarly, the results of Di Luca et al. (2011) suggest that the reliability of proprioceptive information can be lowered by making participants use only their shoulder joint to press down on an object rather than using the wrist and elbow. For such reasons, here we will assume that the reliability of proprioceptive information about force r_{fp} can be either higher or lower than the reliability of tactile information r_{ft}. Fig. 5.4a shows the condition where the final estimate of compliance is obtained using displacement information and an integrated force difference estimate. In such a case where integration of force according to MLE precedes combination of position and force signal (integration before combination), the final estimate of compliance can be expressed as such:

$$C_{i\_c} = p/f = p/(w\, f_p + (1\text{-}w)\, f_t) \qquad (5.3),$$

where the weight is calculated according to reliability of the force estimate r_{fp} and r_{ft} as such

$$w = r_{fp} / (r_{fp}+r_{ft}) \qquad (5.4).$$

The alternative (combination before integration) illustrated in Fig. 5.4b shows two estimates of compliance obtained using a unique proprioceptive position signal. The two compliance estimates are redundant and they can be integrated according to their reliability r_{cp} and r_{ct}. The result can thus be expressed as:

$$C_{c\_i} = w\, C_p + (1\text{-}w)\, C_t = w\, (p/f_p) + (1\text{-}w)\, (p/f_t) \qquad \text{(Eqn. 5.5)}.$$

Here the weight w should be calculated according to reliability of the compliance estimates, not the reliability of the force estimates as in the previous case. The difference with this procedure is that both estimates of compliance use the same position information. In this case, the noise affecting the two compliance estimates is not independent. Thus, instead of using the simple formula

$$w = r_{cp} / (r_{cp}+r_{ct}) \qquad (5.6),$$

the weighting scheme should be modified to account for the correlation between noises. The weight needs to be calculated according to the following formula (see Oruç et al. 2003):

$$w = r_{cp}' / (r_{cp}' + r_{ct}') \quad (5.7).$$

Here r' is the corrected reliability that accounts for the correlation ρ between the noise in the two compliance estimates and it corresponds to

$$r_{cp}' = r_{cp} - \rho\, r_{cp} r_{ct} \quad (5.8)$$

and

$$r_{ct}' = r_{ct} - \rho\, r_{cp} r_{ct} \quad (5.9).$$

We performed numerical simulations to evaluate the reliability of the compliance estimate obtained following Eqn. 5.3 and Eqn. 5.5. The frequency of the obtained difference in reliability for the two methods is summarized in Fig 5.5a. Results indicate that there is a demarcated difference between the two methods. At intermediate compliance levels, estimating obtained by firstly integrating redundant force information and only then combining position and force could have higher reliability than the one obtained by firstly combining and then integrating. This is most likely because force estimates obtained with the two modalities have been simulated to differ only moderately. In such cases, the correlation of the noise in the compliance estimates makes the weighting deviate from optimality (Fig. 5.4b) decreasing the reliability of the posterior (Oruç et al., 2003).

3.2 Proprioception and Vision

Let us now analyze the case of a compliance estimate obtained with redundant position information provided by vision and proprioception (p_v. and p_p, respectively). Again, this is a case where tactile information does not contribute to compliance estimate directly because contact points are rigid. In such a condition, the visual sense is much more precise than proprioception for estimating object indentation. Redundant position information is integrated before combining it with force to obtain a compliance estimate as shown in Fig. 5.4c. The final estimate of compliance can be expressed by:

$$C_{i\_c} = (w\, p_v + (1-w)\, p_p) / f \quad (5.10),$$

where the weight w is calculated according to reliability of the two position estimates r_{pv} and r_{pp} and thus it is expected to be high for vision:

$$w = r_{pv} / (r_{pv} + r_{pp}) \quad (5.11).$$

On the other hand, integrating redundant estimates of compliance after position and force have been combined as shown in Fig. 5.4c would lead to a compliant estimate equal to

$$C_{c\_i} = w\ C_v + (1\text{-}w)\ C_p = w\ (p_v/f) + (1\text{-}w)\ (p_p/f) \qquad (5.12),$$

but again the weight is not calculated according to the simple reliabilities

$$w = r_{cv} / (r_{cv}+r_{cp}) \qquad (5.13),$$

but instead is computed according to the corrected reliability values shown above, which consider that noise is not independent

$$w = r_{cv}' / (r_{cv}'+r_{cp}') \qquad (5.14),$$

thereby taking into account the correlation between the noise of the estimates due to the use of the same force signal. The results of the numerical simulation displayed in Fig. 5.5b show the difference in reliability between compliance obtained following Eqn. 5.10 and Eqn. 5.12. By accounting for the correlation between variances of the estimates, there are small differences in the reliability attainable with the two methods, with a modest improvement in performance by first combining information in separate compliance estimate, and then integrating them. This is the result of improvement in precision due to the presence of two estimates of compliance. Such improvement is lost when integrating position information because of the extreme weight that should be assigned to the visual estimate of position.

Results of the simulation suggest that integrating force information to obtain a unique force estimate to be combined with position information might lead to a more precise overall compliance estimate. On the other hand, integrating position information available through proprioception and vision leads instead to a slightly worse final estimate of compliance than first obtaining two separate estimates of compliance to be integrated.

The finding that the brain is better off by combining force and position signals when visual and proprioceptive modalities are available is consistent with the finding of Kuschel et al. (2010). The researchers found that compliance reliability is very low when only visual position information and haptic force information are available. Such combination of sensory information is not sufficient to estimate compliance if the brain first computes a compliance estimate and then integrates the various estimates available (see Chapter 2 for an in depth discussion). It would be interesting to perform a similar test with touch and proprioceptive information (force information available through touch, but not through proprioception), as for this case the best performance can be achieved by first integrating force signals and only then combine with position (Fig. 5.5a).

To summarize, integrating compliance information, rather than position or force information, should depend on the correlation between the noise sources affecting the sensory signals involved in the process. One could ask whether such a difference is also present when multiple spatially-separated sensors are available (i.e., when interaction with an object is performed with multiple fingers). Such a case will be considered in the next section.

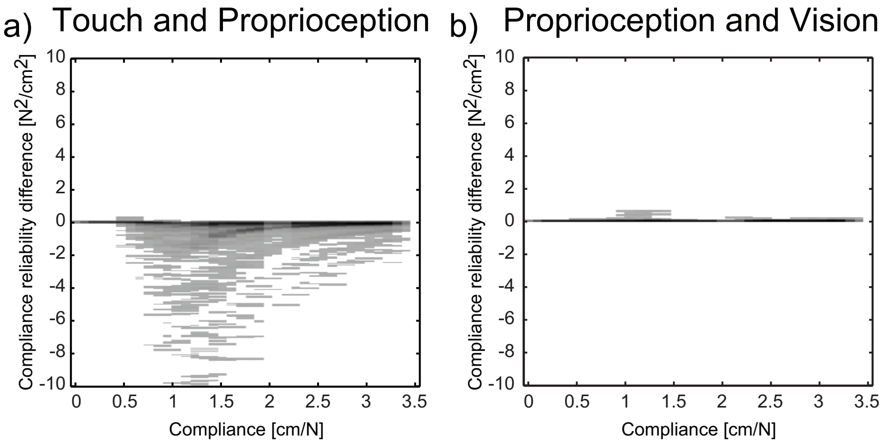

Fig. 5.5. Numerical simulation comparing the reliability of the compliance estimate obtained in 10000 samples with the two methods in the two scenarios described in Fig. 5.4. Weber Fraction values between 0.1 and 0.5 have been employed for each of the signals involved. Darker color indicates more frequent results. Positive values of the ordinate indicate a more reliable result with a scheme that first combines the signals and then integrates the redundant compliance estimates. If the two methods were mathematically equivalent, we would expect a flat line at 0. a) Force can be obtained through tactile and proprioceptive modalities, while position is estimated through proprioception. The shaded area below the zero line indicates cases where integration before combination (Eqn. 5.3) can produce higher overall reliability than what is obtained with combination before integration (Eqn. 5.5). b) Force is obtained through proprioception and position is obtained redundantly through vision and proprioception. There are very few cases where the reliability is different for the two types of estimates, where the combination before integration (Eqn. 5.10) leads to a better result than Eqn. 5.12.

4 Multiple Contact Points

Our perception of the compliance seems to depend on the way we interact with of objects (Di Luca, 2011; Kaim & Drewing, 2009). The overall judgment of the quality of the fruit depends on the active exploration and the pattern of movement of the fingers. So how much movement do we plan to apply to begin with? Chapter 6 shows how, for example, the amount of force depends on the difficulty of the anticipated discrimination, and on the compliance of the object that is expected, as using the appropriate force would increase the discrimination sensitivity. Moreover, with very compliant objects better performance can be achieved using 3 fingers rather than only 1 (see Chapter 6). Another question that has not been addressed is whether participants *choose* whether to use one or multiple fingers in the interaction and when do they decide to switch. Finally, it is not clear whether using multiple fingers in contact with compliant objects leads to a specialized role of each finger. Some contacts might be used predominantly for stabilisation while other might be used for active exploration. In extreme cases when one presses down on an object with just one finger there would be only one contact point that provides information about compliance. Another way to assess the softness of the object would be to grasp it with two fingers in a "pinch" grip. Employing such a precision grasp could gather information about material properties from both contact points. Each contact point is one sensor that collects sensory information that needs to be integrated appropriately. But the information they collect could vary for example if the grip is horizontal or vertical.

There are also cases when compliance information is obtained with a full hand grasp, or even two hands. Here, the information from each finger and contact point can be used as a separate estimate of the overall compliance of the object. In such cases of multiple contacts, we need to consider that compliance is not constant across the object, i.e., fruit might have some areas where the material is soft because it starts to rot. Thus, with multiple contact points, the spatially separated estimates need to be integrated into one percept of the object compliance, assuming an object with compliance properties that are similar throughout its surface.

Many studies investigating compliance perception with objects having rigid surfaces mostly analyzed single contact-point interactions (Srinivasan & LaMotte, 1995; Di Luca et al., 2011). The few cases that analyzed more than one contact point, as it happens in the case of object holding (Chen & Srinivasan, 1998) or grasping (Roland & Ladegaard-Pedersen, 1977; Tan, Durlach, Beauregard, & Srinivasan, 1995; Freyberger & Färber, 2006; Kuschel et al., 2011; Di Luca, 2011), considered the sensory information as the sum of the two forces and deformations. If the object is squeezed between fingers, i.e., contact points are producing forces in opposite directions, one could assume that forces and deformations at the two ends should be simply summed up. Such assumption, however, is not always fulfilled, for example, as in the case where objects are pressed down with two fingers. Moreover, by considering that with pinch grasps it is possible to generate independent finger movements (Schieber, 1996; Smeets & Brenner, 2001), it becomes apparent that the fingers can act as separate sensors to collect information about resistive force and deformation. A simple sum of the two sources of information would not be statistically optimal in all situations (it would not lead to the most precise estimate of compliance). To obtain an optimal estimate of compliance from two sensory sources, each source should be weighted according to its reliability, consistent with Bayesian inference (Knill & Richards, 1996). Indeed it has been shown that in many instances the perceptual system is, in principle, capable of obtaining close to optimal performance in compliance judgments obtained with multiple contact points (Di Luca, 2011; Plaisier et al., 2012). The weighting of each contact estimate changes as a function of the exploratory movement reflecting the reliability of the estimate (Di Luca, 2011). This weighting scheme is consistent with the relative reliability of compliance estimates when the object is composed of a uniform material. Performance obtained in this case is close to optimal. Others have found a close to optimal performance in the case of consistent sources of compliance information (Kuschel et al., 2010; Di Luca et al., 2010).

The perceptual system should only integrate information that comes from the same distal source and should behave in a robust way otherwise. That is, when conflicts or evidence about the origin of sensory information indicate the presence of separate sources, integration should simply not occur. For example, weighted averaging should not occur mandatorily when large conflicts are present between different sources of information (i.e., van Ee, van Dam, & Erkelens, 2002). Di Luca (2011) also finds that perceived compliance is mostly dictated by one source of information when the conflict is large, however the source chosen to drive the percept is not the one that provides more reliable information. Different results show, in fact, that there is a lawful relation between reliability and overall compliance (Jones & Hunter, 1990; Tan et al., 1995). If only compliance judgment reliability was at stake, there should be a consistent weighting according to the reliability difference. Instead, results by Di Luca (2011) suggest that when using a precision grip, it is the information coming from the finger that moves the most that drives the percept. Work in other domains also showed that the source of sensory information chosen to drive the percept was not necessarily the more reliable one (Girshick & Banks, 2009; Gori et al., 2008). Overall these results are in line with a scheme of integration that can lead to near-optimal perception, but beyond the limit of fusion progressively increases the weighting of one source of information. It does not, however do this based on reliability alone.

5 Temporal Aspects of Softness Perception

Whether the computation proceeds by first integrating and then combining, or by first combining and then integrating; either way, force and position information are necessary to obtain a softness percept of objects with rigid surfaces (Fig. 5.1a). Several studies investigating compliance perception of objects with rigid surfaces make two major assumptions:

- First, that the objects to be manipulated are in “unloaded” resting state, so that even infinitesimal forces can create an indentation albeit infinitesimal. See Tan et al. (1995) for an exception.
- The second premise is that one of the two indentation positions considered for the estimate of compliance is the resting state of the object (the non-indented position) and thus, only one position is analysed (often the one at the maximum indentation).

These simplifications offer the advantage of reducing the problem of estimating compliance to the estimate of a single force and a single position, seemingly obtained at the point where force reaches the maximum (see Tan et al., 1995 for an investigation). These premises, however, are not always fulfilled and create computational problems when dealing with non-Hookean materials. For example, if force and position are related to the starting point of indentation (i.e., the point at which the first contact is made with the objects surface), then the non-indented position of the object should have a paramount influence on perception (see Nisky et al., 2008). Pressman et al. (2011) indeed showed that the unnoticed movement of the object could cause a misperception of compliance that can be ascribed to the observer using previous estimates of the non-indented position. This information would act as a prior for position to be integrated at the force/position level. Interestingly, small force gradients of very compliant objects could lead to biases in the estimation of the perceived location of the non-indented position. Namely, the perceived location of the non-indented surface is biased inwards due to sensory thresholds for force (Chapter 7 includes a detailed analysis of the influence of the perceived location of the boundary of a force field). In other words, if compliance is estimated only through the haptic modality, and information using the non-indented position of the object is used as a reference, perceived softness with very compliant objects should be lower than veridical (i.e., very compliant objects should be perceived less compliant than they are). Chapter 9 analyses the case of the interaction with haptic interfaces where force rendering is delayed with respect to position. With a delay in the force generated by the device, the perceived position of the non-indented object should be biased to a position further inside the object (Fig. 5.3d). Consequently, the slope used to estimate compliance (Fig. 5.1a) should be higher, leading to a less compliant percept. Interestingly, there is an indication that with haptic-only interaction, a delay in the rendering of force leads to underestimation of compliance – objects are perceived to be harder when a force delay is present (Pressman et al., 2007). On the contrary, adding visual information about the indentation, in particular about the position of the non-indented surface of the object, reverses the effect making force-delayed objects appear more compliant than non-delayed ones (Ohnishi & Mochizuki, 2007; Di Luca et al., 2011).

It is important to note that integration of compliance information allows the use of multiple estimates, which has the advantage of improving reliability. A positive correlation between the noises corrupting the sensory information, however, leads to a reduction of the overall reliability of the final, integrated estimate of compliance (i.e. Oruç et al., 2003). Moreover, integrating redundant compliance information, rather than redundant position or force, allows the brain to use prior knowledge about compliance in the final estimate. Prior knowledge about compliance is independent of the actual force and position values. It is known that knowledge about the regularities of the world can help to reduce uncertainty and ambiguity in perception by creating an *a priori* expectation of what the most probable state of the world is (Knill & Richards, 1996). The knowledge acquired from past experience interacting with the object, e.g., about its compliance, can be expressed using a probability distribution across all possible compliances. *A priori* information has a strong effect on perception in everyday situations. Namely, priors actually influence many of the properties of the perceptual world and have a substantial impact on the dynamics of sensory processing. In particular, priors act as a predictor for the state of the world for which motor actions are initially planned. When approaching an object in order to sense its softness, the brain needs to have an approximate idea of what the softness of said object is, so that actions after contact are executed accordingly. Actions are planned on the basis of *a priori* expectations about the material properties. Such internal models based on previous interactions (Kawato et al., 1987) are triggered from the view of the object before contact. In other words, recognizing an object will activate the internal model and create an expectation of what it will feel like when in contact. This expectation should be used for both perception and action. We discussed the influence of priors on perception, but what happens on action could be resumed by considering the example situation where we expect the object to be very soft. Here we should NOT approach the object by applying a large force (i.e., a strong grip force), as this will compress the object entirely. We should instead approach the object using a low force to perturb it, so as to increase the precision of force and position signals used to estimate compliance. In other words, the best strategy to improve the precision of a perceptual judgment is to base the movement on prior knowledge about how the object will react, so to detect any discrepancy in compliance from the initial guess.

Another problem to overcome is that there is a delay between the motor command and the physical changes in the environment that will generate sensory information. This constitutes a computational problem that could be solved using a forward model about the motor action (Miall & Wolpert 1996). This model should also be updated if it is not accurate when new sensory information comes about. In other words, as sensory signals become available after contact, the influence of the internal model of the material (prior knowledge) should decrease and discrepancies between the forward models and the actual state of the world should be used to update the estimate, as well as to change the movement parameters.

What complicates things even further is that the estimate of compliance of objects having rigid surfaces requires an active indentation of the object (i.e., a *change* in the global deformation of the object through the application of a different force). Because of this, the estimate of compliance undergoes continuous updating as more sensory information becomes available over time. The Kalman filter is a statistically optimal method that can update estimates over time. Human performance has been hypothesized to employ computational mechanisms similar to the Kalman filter in sensorimotor integration (Wolpert, 1997) and in perception (Rao, 1999), but it is not clear whether such a statistically optimal update happens to the internal model of compliance.

Notably, as the estimate of compliance is dynamic, it is not possible to know experimentally what the perception of softness is at any *one* moment during the interaction. Experiments have usually employed the task of judging the object compliance after the interaction has been completed. Similarly, the concept of minimizing uncertainty over the course of the exploration cannot be extended directly to every time point during the interaction. The series of movements needed to first reach, grasp and then probe a deformable object involves specific cost functions, and each has an accumulation of sensory information with a different goal. For example, parts of the task, like reaching and grasping, might involve minimizing energy consumption or execution time, while other components of the task could consider the chances of hitting other objects. In such cases, looking at overall performance and characterising it as noise reduction in the response to an experimental task might not be the best way. Information accumulation could be better characterized with an analysis of information available and movement performed at different time points over the course of the interaction (Fig. 5.6). For example, when estimating the size of a bar using touch and vision we accumulate tactile information over the first second, while the time course is much faster for vision (see Ernst, 2001). Thus, longer trial durations would result in a more precise size estimate through the haptic information, but no comparable increase should be present for visual information. In the same way, during interactions with compliant objects the quality of haptic sensory information can increase with more prolonged interactions. The situation is further complicated as compliance perception requires movement (i.e., information is gathered through motion, whether this is planned and self-generated, or not).

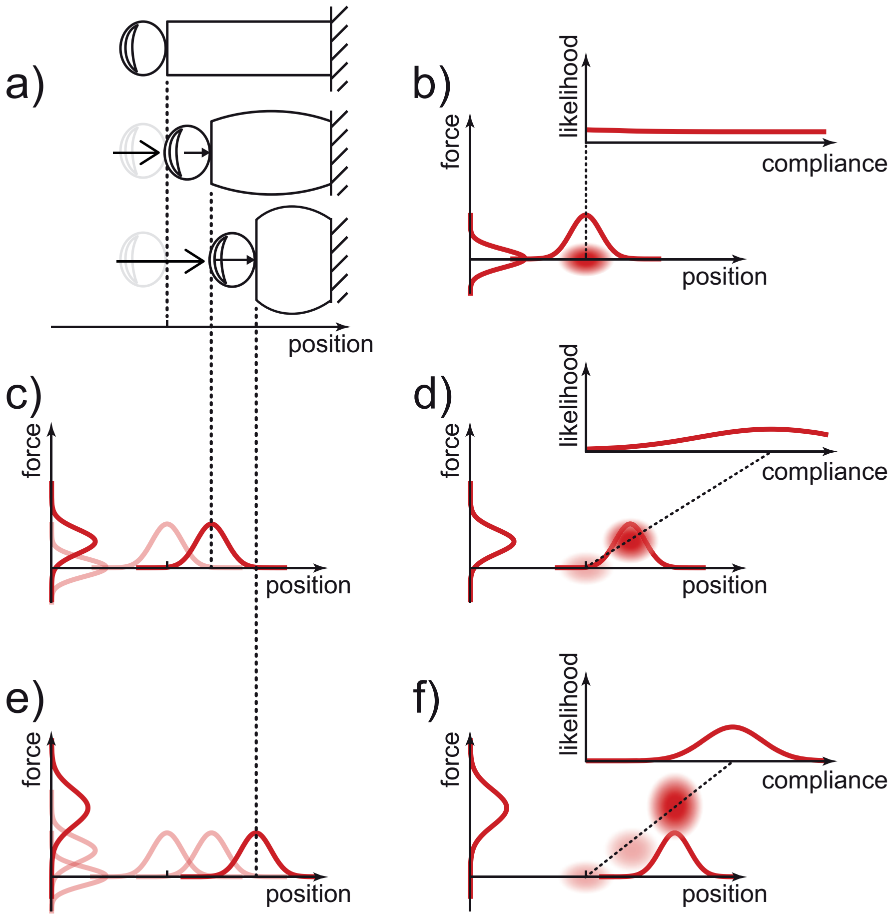

Fig. 5.6. Representation of the interaction with a soft object over time. a) force and position sampling distributions. b) the reliability of the compliance estimate increases as more sampling points become available.

Di Luca et al. (2011) have shown that the phases of the indenting action (loading the object with a force and unloading it) are differently informative regarding compliance. Namely, unloading actions are not as informative about compliance as loading ones. You could demonstrate this for yourself by comparing the availability and strength of a softness percept while pushing against an object and releasing the applied force. This means that when an object is pushed against, there can be many sensory signals on which the movement planning can be optimized upon (amount of deformation over time with correspondingly resistive force values, vibratory and auditory patterns, skin pressure distribution, stretching dynamics, etc.). Movements can thus change after each repetition so to increase the reliability of the final estimated property.

Integration of multiple sources of sensory information has been shown to be close to optimal when stimuli are presented simultaneously, they have short duration and the specified property doesn't change over time (see Ernst & Banks, 2002; Ernst & Bülthoff, 2004). This leads to a normative solution for any discrepancies and to an increase in the precision of perceptual estimates. The question here is whether such an increase in reliability happens also when the stimulation is accumulated over time (i.e. sequentially, Fig. 5.6) and how perceptual performance is affected. In compliance perception we expect integration to lead to some increase in performance, but not as much as with perceptual situations with multiple independent estimates because of the high temporal correlation at subsequent time points (Oruç et al., 2003, Juni et al., 2012).

6 Conclusions

In this chapter we presented some of the computational properties of the mechanisms involved in compliance perception. In particular, we proposed a model of human perception of compliance based on Bayesian Decision Theory. The model indicates how complementary and redundant information should be treated. First of all, the brain should make use of sensory information obtained during the interaction, as well as knowledge about the object properties obtained from previous interactions. Moreover, the model shows how sensory information obtained from different modalities should be treated to obtain a compliance estimate. In some conditions it might be profitable to obtain redundant estimates of compliance, whereas in other situations precision would be improved by using a unique position and force estimates from all signals available. The adaptability of the process leading to a compliance estimate is also underlined by considering the dynamic nature of information about compliance, how the information is accumulated over time, and how it is combined across multiple contact points. Results of this analysis suggest that one reason that compliance is normally perceived through multiple sensory signals, across several sense modalities, and over time is that such a range of signals leads to an improvement of the quality of the final compliance estimate. Such improvement is most important if direct contact with the object does not provide direct information about its softness, as is the case when interacting with objects having rigid surfaces.

7 Acknowledgment

The authors are grateful to Markus Rank and Darren Rhodes for help in preparing the manuscript.

8 References

Backus, B. T., & Banks, M. S. (1999). Estimator reliability and distance scaling in stereoscopic slant perception. Perception, 28, 217–242.
Bergmann-Tiest, W. M., & Kappers, A. M. L. (2009). Cues for Haptic Perception of Compliance. IEEE Transactions on Haptics, 2(4), 189–199.
Bicchi, A., Scilingo, E. P., & De Rossi, D. (2000). Haptic discrimination of softness in teleoperation: the role of the contact area spread rate. Robotics and Automation, IEEE Transactions on, 16(5), 496–504.
Bresciani, J.-P., M. O. Ernst, M. O., K. Drewing, K., G. Bouyer, G., V. Maury, V., & and A. Kheddar, A. (2005). Feeling what you hear: auditory signals can modulate tactile taps perception. Experimental Brain Research , 162(2), 172-180.
Cellini, C., Kaim, L., & Drewing, K. (2013). Visual and haptic integration in the estimation of softness of deformable objects. I-Perception, 4(8), 516–531.
Chen, J. S., & Srinivasan, M. A. (1998). Human haptic interaction with soft objects: Discriminability, force control, and contact visualization. MIT, RLE Technical report 619 (pp. 1–208).
Chen, J.-S., Srinivasan, M. A., Massachusetts Institute of Technology. Research Laboratory of Electronics, Massachusetts Institute of Technology. Laboratory of Human and Machine Haptics. (1998). Human Haptic Interaction with Soft Objects.
Di Luca, M. (2011). Perceived compliance in a pinch. Vision Research, 51(8), 961–967.
Di Luca, M., Knorlein, B., Ernst, M. O., & Harders, M. (2011). Effects of visual–haptic asynchronies and loading–unloading movements on compliance perception. Brain Research Bulletin, 85(5), 245–259.
Duda, R. O., Hart, P. E., & Stork, D. G. (2000). Pattern classification (2nd ed.). New York: John Wiley & Sons.
Ernst M.O. (2001). Psychophysikalische Untersuchungen zur visuomotorischen Integration beim Menschen: visuelle und haptische Wahrnehmung virtueller und realer Objekte, Eberhard-Karls-Universität Tübingen,
Ernst, M. O. (2006). A Bayesian view on multimodal cue integration. In Knoblich, G., Thornton, I. M., Grosjean, M., & Shiffrar, M. (Eds.) Human Body Perception From the Inside Out, 105–131, NY: Oxford University Press.
Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. Nature, 415(6870), 429–33.
Ernst, M. O., & Bülthoff, H. H. (2004). Merging the senses into a robust percept. Trends in Cognitive Sciences.
Freyberger, F. K. B., & Färber, B. (2006). Compliance discrimination of deformable objects by squeezing with one and two fingers. Proceedings of EuroHaptics, 06, 271–276.
Girshick, A. R., & Banks, M. S. (2009). Probabilistic combination of slant information: Weighted averaging and robustness as optimal percepts. Journal of Vision, 9(9), 1–20.
Gori, M., Del Viva, M., Sandini, G., & Burr, D. C. (2008). Young Children Do Not Integrate Visual and Haptic Form Information. Current Biology, 18(9), 694–698.
Helbig, H. B., & Ernst, M. O. (2007). Knowledge about a common source can promote visual haptic integration. Perception, 36, 1523-1533.
Jones, L. A., & Hunter, I. W. (1990). A perceptual analysis of stiffness. Experimental Brain Research, 79, 150–156.
Juni, M. Z., Gureckis, T. M., & Maloney, L. T. (2012). Effective integration of serially presented stochastic cues. Journal of Vision, 12(8), 12–12.
Kawato, M., Furukawa, K., & Suzuki, R. (1987). A hierarchical neural-network model for control and learning of voluntary movement. Biological Cybernetics, 57(3), 169–185.
Kaim, L. & Drewing, K. (2009). Finger force of exploratory movements is adapted to the compliance of deformable objects. In Proceedings, World Haptics 2009, Third Joint EuroHaptics Conference and Symposium on Haptic Interfaces for Virtual Environment and

Teleoperator Systems (pp. 565- 569). The Institute of Electrical and Electronics Engineers (IEEE): Piscataway, NJ.
Knill, D. C., & Richards, W. (1996). Perception as Bayesian Inference Cambridge University Press.
Kuschel, M., Di Luca, M., Buss, M., & Klatzky, R. L. (2010). Combination and Integration in the Perception of Visual-Haptic Compliance Information. IEEE Transactions on Haptics, 3(4), 234–244.
Mamassian, P., Landy, M. S., Maloney, L. T. (2002). Bayesian modelling of visual perception. In Rao, R., Olshausen, B. A., & Lewicki, M. S. (Eds.) Probabilistic Models of the Brain: Perception and Neural Function, 13–36.
Miall, R. C., & Wolpert, D. M. (1996). Forward models for physiological motor control. Neural Networks, 9(8), 1265–1279.
Nisky, I., & Mussa-Ivaldi, F. A. (2008). A regression and boundary-crossing-based model for the perception of delayed stiffness. IEEE Transactions on Haptics, 1(2), 73-83.
Ohnishi, H., & Mochizuki, K. (2007). Effect of Delay of Feedback Force on Perception of Elastic Force: A Psychophysical Approach. IEICE TRANSACTIONS on Communications, E90-B(1), 12–20.
Oruç, İ., Maloney, L. T., & Landy, M. S. (2003). Weighted linear cue combination with possibly correlated error. Vision Research, 43(23), 2451–2468.
Parise, C. V., Spence, C., & Ernst, M. O. (2011). When Correlation Implies Causation in Multisensory Integration. Current Biology, 1–4.
Plaisier, M. A., & Ernst, M. O. (2012). Two hands perceive better than one. Haptics: Perception, Devices, Mobility, and Communication, Lecture Notes in Computer Science, 7283, 127–132.
Pressman, A., Karniel, A., & Mussa-Ivaldi, F. A. (2011). How Soft Is That Pillow? The Perceptual Localization of the Hand and the Haptic Assessment of Contact Rigidity. Journal of Neuroscience, 31(17), 6595–6604.
Roland, P. E., & Ladegaard-Pedersen, H. (1977). A quantitative analysis of sensations of tensions and of kinaesthesia in man. Brain, 100, 671–692
Rao, R. P. (1999). An optimal estimation approach to visual perception and learning. Vision Research, 39(11), 1963–1989.
Roskies, A.L. (1999), The binding problem, Neuron, 24, 7-9
Schieber, M. H. (1996). Individuated finger movements: Rejecting the labeled-line hypothesis. In A. M. Wing, P. Haggard, & J. R. Flanagan (Eds.), Hand and brain (pp. 81–98). New York: Academic Press.
Scilingo, E. P., Bianchi, M., Grioli, G., & Bicchi, A. (2010). Rendering softness: Integration of kinesthetic and cutaneous information in a haptic device. IEEE Transactions on Haptics, 3(2), 109–118.
Smeets, J. B. J., & Brenner, E. (2001). Independent movements of the digits in grasping. Experimental Brain Research, 139, 92–100.
Srinivasan, M. A., & LaMotte, R. H. (1995). Tactual discrimination of softness. Journal of Neurophysiology, 73(1), 88–101.
Tan, H. Z., Durlach, N. I., Beauregard, G. L., & Srinivasan, M. A. (1995). Manual discrimination of compliance using active pinch grasp: The roles of force and work cues. Perception & Psychophysics, 57(4), 495–510.
Thurlow WR, Jack CE. (1973). Certain determinants of the "ventriloquist effect". Perceptual and Motor Skills, 36, 1171-1184
van Ee, R., van Dam, L. C. J., & Erkelens, C. J. (2002). Bi-stability in perceived slant when binocular disparity and monocular perspective specify different slants. Journal of Vision, 2(9), 2–2.
Visell, Y., Giordano, B. L., Millet, G., & Cooperstock, J. R. (2011). Vibration influences haptic perception of surface compliance during walking. PloS One, 6(3), e17697.
Witkin, H. A., Wapner, S., & Leventhal, T. J., (1952). Sound localization with conflicting visual and auditory cues, Journal of Experimental Psychology, 43, 58-67.

Wolpert, D. M. (1997). Computational approaches to motor control. Trends in Cognitive Sciences, 1(6), 209–216.