Adaptation effects in grasping the Müller-Lyer illusion

Karl K. Kopiske [a], Evan Cesanek [b], Carlo Campagnoli [b, c], Fulvio Domini [a, b]

a: Center for Neuroscience and Cognitive Systems @UniTn, Istituto Italiano di Tecnologia, Corso Bettini 31, 38068 Rovereto, Italy.

b: Department of Cognitive, Linguistic and Psychological Sciences, Brown University, Providence, RI 02912, USA.

c: Department of Psychology, Princeton University, Princeton, NJ 08544, USA.

Corresponding author:

Karl K. Kopiske [a]

email:

karl.kopiske@iit.it

address:

Center for Neuroscience and Cognitive Systems@UniTn,

Istituto Italiano di Tecnologia,

Corso Bettini, 31,

38068 Rovereto, TN,

Italy.

Running Head: Adaptation to two illusion displays in grasping

Manuscript overall word count: 9596 (main text only: 6526).

Abstract

Recent results have shown that effects of pictorial illusions in grasping may decrease over the course of an experiment. This can be explained as an effect of sensorimotor learning if we consider a pictorial size illusion as simply a perturbation of visually perceived size. However, some studies have reported very constant illusion effects over trials. In the present paper, we apply an error-correction model of adaptation to experimental data of N = 40 participants grasping the Müller-Lyer illusion. Specifically, participants grasped targets embedded in incremental and decremental Müller-Lyer illusion displays in (1) the same block in pseudo-randomized order, and (2) separate blocks of only one type of illusion each. Consistent with predictions of our model, we found an effect of interference between the two types when they were presented intermixed, explaining why adaptation rates may vary depending on the experimental design. We also systematically varied the number of object sizes per block, which turned out to have no effect on the rate of adaptation. This was also in accordance with our model. We discuss implications for the illusion literature, and lay out how error-correction models can explain perception-action dissociations in some, but not all grasping-of-illusion paradigms in a parsimonious and plausible way, without assuming different illusion effects.

Keywords: Müller-Lyer illusion, grasping, adaptation, error correction.

THIS PAPER HAS BEEN ACCEPTED FOR PUBLICATION IN *VISION RESEARCH*!

For a full and up-to-date version, please visit http://doi.org/10.1016/j.visres.2017.05.004

1. Introduction

1.1 Sensorimotor adaptation and visual illusions

When performing repeated motions towards a seen object, humans will rapidly become more adept at this task, a fact already described over a hundred years ago (von Helmholtz, 1867; Woodworth, 1899). It has been proposed that what is learned is an efficient transformation between the visual input and the motor action required for the task, called a visuomotor mapping (Soechting & Flanders, 1989). This holds true under natural conditions, but also when systematic distortions are introduced, either of the visual input (e.g. through mirror setups, Säfström & Edin, 2004) or the motor output (Thoroughman & Shadmehr, 2000). Under such experimental perturbations, participants rapidly adjust their visuomotor mapping to the task demands. Specifically, this continuous adaptation process is believed to be driven by sensory error signals, which can be defined as the difference between a predicted sensory outcome and the observed outcome (Cheng & Sabes, 2006; Shadmehr, Smith, & Krakauer, 2010).

The prevalent use of visual perturbations in sensorimotor adaptation research creates a natural intersection with another area of vision research. For the past 20+ years, researchers have been investigating whether vision for conscious perception is processed in a fundamentally different way to vision used for motor actions (Milner & Goodale, 1995, 2006). Much of the evidence in favour of this theory has come from observations of neuropsychological patients (prominently, Goodale et al., 1994; Goodale, Milner, Jakobson, & Carey, 1991), while evidence in healthy participants has relied strongly on whether and to what degree skilled movements are affected by visual illusions (Aglioti, DeSouza, & Goodale, 1995). This long-standing debate is still ongoing, as some authors (e.g., Goodale, 2014; Westwood & Goodale, 2011) emphasise that there is a large body of literature where grasping has been dissociated from perception in visual illusions, while others argue that when motor and perceptual tasks are well matched, no dissociation is detectable (Franz, Gegenfurtner, Bülthoff, & Fahle, 2000; Kopiske, Bruno, Hesse, Schenk, & Franz, 2016; Schenk, Franz, & Bruno, 2011). However, see also Smeets and Brenner (2006), for a different interpretation that proposes that MGA is not an ideal variable to measure size processing in grasping tasks. While measures other than MGA such as grasp position (Smeets & Brenner, 1999, 2006) or grip force (Hesse, Miller, & Buckingham, 2016; Jackson & Shaw, 2000) have certain advantages, our investigation focusses on the maximum grip aperture (MGA) due to its popularity in the perception-action literature and its well-investigated relationship with object size (Jeannerod, 1984, 1986; Smeets & Brenner, 1999).

Several studies have investigated motor adaptation in some illusions (e.g., Buckingham & Goodale, 2010; Flanagan & Beltzner, 2000; Glover & Dixon, 2001). However, the specific question of whether the proposed smaller sensitivity of grasping to visual size illusions (Goodale, 2014) may be brought about by sensorimotor adaptation (i.e., an illusion effect decreasing over repeated grasping trials, such that the mean value becomes smaller than the initial effect), as well as mechanisms involved have until recently been barely addressed. Whitwell, Buckingham, Enns, Chouinard, and Goodale (2016) reported that the effect of the Ponzo illusion on the MGA decreases substantially over repeated grasping trials performed with full vision of the hand and the object. Similar results were found by Cesanek, Campagnoli, and Domini (2016), also under full-vision conditions, emphasising that illusion effects on grasping may indeed be reduced with practice. If we consider a visual size illusion as in essence a distortion of perceived size, this is exactly what we should expect according to the logic of sensorimotor adaptation. Over many trials of repeatedly grasping objects within the same illusory context, participants may use feedback to learn to adapt their grip scaling to the distortion, much like in visual size perturbation paradigms testing grasp adaptation (e.g., Säfström & Edin, 2004). We should mention that this interpretation is not undisputed, with Whitwell and colleagues (2016) interpreting their results as learning to ignore parts of the illusion that might have been perceived as obstacles, rather than learning a mapping between perceived and grasped size.

Another important point is that while a number of studies found learning effects in grasping illusions, this is not a ubiquitous finding. For example, a study using the Müller-Lyer illusion and the parallel-lines illusion (Franz, Fahle, Bülthoff, & Gegenfurtner, 2001) analysed illusion effects over trials using linear regression, finding no illusion-effect decrease over the course of an experiment. Similarly, reanalysing the open-access data from a recent large-sample study of the Ebbinghaus illusion (Kopiske et al., 2016) also showed no signs of a decreasing illusion effect, which raises the question what the difference may be between studies where a decrease occurs (e.g., Cesanek et al., 2016; Whitwell et al., 2016) and studies where none occurs (e.g., Franz et al., 2001; Kopiske et al., 2016).

1.2 One illusion or many? Our study

When comparing experimental designs between studies that found decreasing illusion effects – e.g., Whitwell et al.'s (2016) and Cesanek et al.'s (2016) experiments – and those that did not – e.g. Franz and colleagues (2001) and Kopiske and colleagues (2016) – some key factors stand out. While all experiments included illusion configurations with opposite (incremental vs.

decremental) effects, these were always presented in the same position in Franz et al. (2001) and in Kopiske et al. (2016), but always in different positions in Cesanek et al. (2016) and Whitwell et al. (2016). However, effects of adaptation can be strongly reduced by spatial separation between response positions (Ghahramani, Wolpert, & Jordan, 1996), such that interference effects between adaptation to one illusion configuration and adaptation to the opposite illusion configuration may also be strongly reduced (Woolley, Tresilian, Carson, & Riek, 2007). Thus, one explanation is that interference would show strongly in the first case (Franz et al., 2001; Kopiske et al., 2016) but less strongly in the second (Cesanek et al., 2016; Whitwell et al., 2016), making it difficult to compare results from the two types of studies.

To investigate these issues, we conducted an experiment in which participants were asked to grasp the Müller-Lyer illusion (fig. 1). Grasping always occurred in the same location, thus maximizing the chance of observing interference effects. Blocks with both illusion configurations were compared to blocks with only one illusion present, enabling us to compare adaptation rates with and without the possibility of interference. We quantified adaptation and interference within the computational framework of linear state-space models, which are commonly used in adaptation experiments to measure error correction rates (Cheng & Sabes, 2006; Thoroughman & Shadmehr, 2000). Using such a model enabled us to compare error correction parameters between different blocks. We also analysed the interference between multiple illusion conditions by fitting a generalization parameter that describes how errors experienced during movements toward one target type influence movements toward another target type (cf. Ghahramani et al., 1996; Krakauer, Ghez, & Ghilardi, 2005). Since other differences between previous studies include the number of objects used (higher in Franz et al., 2001, and Kopiske et al. 2016, than in

Whitwell et al. 2016 and Cesanek et al. 2016; four and five respectively, vs. two each), we also manipulated the number of objects per block to preclude this as a confounding factor (for an argument why this factor could matter, see Keefe & Watt, 2009), although our model would not predict this manipulation to affect adaptation.

Following these considerations, we hypothesised that participants should quickly be able to learn to adapt their grasp to one, but not multiple opposite illusory size distortions in the same location, and derived two testable hypotheses from this: (1) We should find the illusion effect decreasing over trials, and (2) in our design, the decrease rate should be smaller with multiple different illusion configurations (pseudo-)randomly interleaved per block. These hypotheses were tested in a design using the Müller-Lyer illusion (fig. 1), an illusion with similar target shapes to the Ponzo illusion, and known to induce a strong size illusion effect (Bruno & Franz, 2009). In fact, the review by Bruno and Franz (2009) also found the number of trials to be negatively correlated with the magnitude of the illusion effect, although this is not sufficient to infer some learning mechanism. We tested grasping under three different conditions: (1) Having participants grasp a single illusion configuration per block, with one single object size; (2) having participants grasp a single illusion configuration per block, but with two object sizes; and (3) having participants grasp both illusion configurations in the same block.

2. Methods

2.1 Sample and setup

We tested a sample of N=40[1] participants (between 18 and 32 years old, mean age = 20.4, right-handed by self-report), recruited at Brown University. Participants completed a simple reach-to-grasp task that was approved by the Brown University Institutional Review Board. Participants gave written, informed consent, received course credit or $8 per hour compensation, and their rights were protected according to the 1964 Declaration of Helsinki.

During the experiment, participants were seated on a height-adjustable stool. In front of them was a table with an armrest and a cardboard grip marking the starting position. Small posts with three infrared diodes each were attached to the thumb and index finger of the participants to track their motion in space at 85 Hz, using an Optotrak Certus motion tracking system (Northern Digital, Waterloo, Canada). Before the experiment, a calibration procedure was performed where participants pinched a predetermined position with their right thumb and index finger to enable us to calculate the position of the respective finger tips relative to the markers on the posts. The

[1] A formal power analysis was not conducted to compute the required sample size, since no single test statistic was decisive for our study and our main focus was on fitting the model. A total of 47 participants were tested, with data from 7 being removed from analysis due to incomplete data. Since adaptation over trials was one of our main concerns, invalid trials were only repeated a maximum of two times, after which the trial was recorded as invalid. If in any bin (see section 2.3) there were not enough trials to compute an illusion effect, the participant was marked as having incomplete data.

finger-tip positions were used in all subsequent analyses. During the experiment, participants placed their head on a chin rest and were instructed to keep their gaze on a semi-transparent mirror (with a removable back panel) in front of them, which allowed the simultaneous presentation of real objects (placed on a carousel app. 50 cm in front of the participants) and virtual objects presented on a 19’’ CRT monitor running at 85 Hz in the same location (see figure 1a). The objects to be grasped were cuboids of 40 or 45 mm length and a 13mm*13mm base. The size of the virtual stimulus always corresponded with the size of the physical object, and an infrared diode on the back of each physical object allowed us to ensure perfect alignment of the virtual and physical object. In the illusion trials, fins of 4 mm * 16 mm extruded from the virtual display at an angle of either 45 or 135 degrees relative to the vertical dimension to induce the Müller-Lyer illusion (see figure 1b). The stimuli were presented in a custom C++ program using OpenGL. The back panel of the mirror was inserted after ensuring visual-physical alignment in order to block participants’ vision of the real objects as well as the hand during non-practice trials of the experiment.

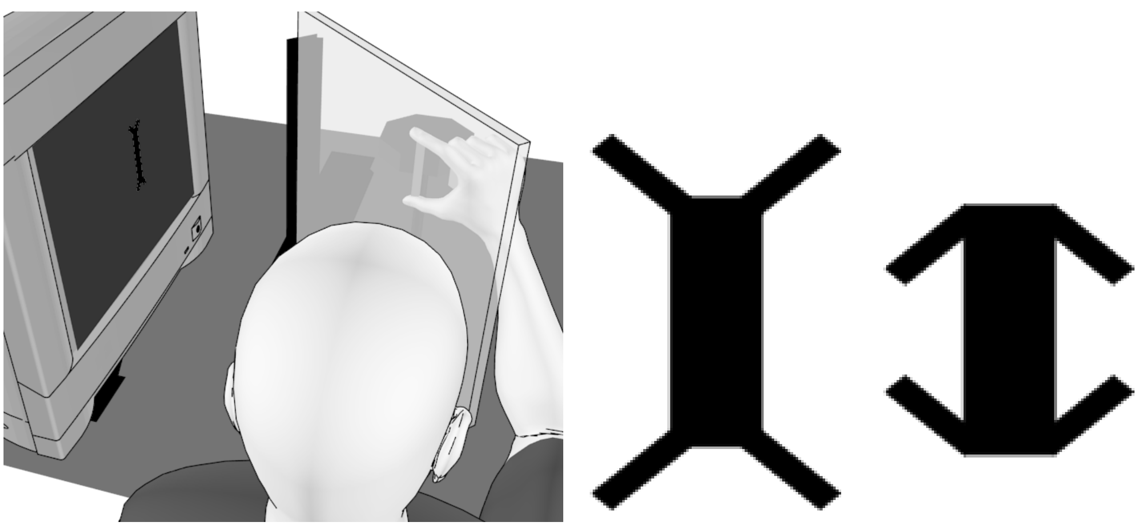

Figure 1: Setup and illusion configurations **a**: A schematic view of the setup with participants looking at a mirror showing the reflection of the screen in the same position as the stimulus behind the mirror. Reach-to-grasp movements were executed from a fixed starting position to a fixed object position. This position was to the bottom right of the object. **b**: Müller-Lyer illusion, as used in the experiment. The virtual stimuli shown in the experiment were red. During trials where no illusion was presented, participants saw a plain red rectangle in the position of the object.

2.2 Procedure

Participants were instructed to keep their right hand in a starting position on the table in front of them, pinching a cardboard grip and waiting for a start beep to start their grasp. A virtual stimulus (see Figure 1; a red rectangle with fins for illusion trials, and without fins in practice, baseline, and washout trials) appeared at the same time as the start beep. After each beep, they reached toward the object in front of them and grasped it along its vertical axis with their thumb and index finger, before returning to the same starting position. Participants did not see their right hand during grasping, but continued to see the virtual object until the physical object was touched. Unlike some previous studies (e.g., Whitwell et al., 2016), we chose not to provide online visual information since we know that this can diminish illusion effects on the MGA, which would reduce the sensitivity of our methods. The visual object in front of them disappeared when a target-reached criterion was met (see section 2.3 for details). The maximum time to complete the grasp was set at 3000 msec (this maximum was virtually never reached; mean movement time was 1019 msec). If this time was exceeded, or participants had started the trial before the start beep, or if the finger markers were not visible to the Optotrak for too large portions of the movement (> 20% of frames during movement), participants heard a high-pitched beep to indicate that the trial was invalid. Invalid trials were repeated a maximum of two times, after which they were marked as not completed (among the final sample, this applied to a total of 27 trials, or 0.2%).

The experiment consisted of five blocks for each participant, each consisting of a baseline phase (12 trials), a test phase (20 or 40 trials), which was when the illusion displays were presented, and a washout phase (12 trials): Two *one-size, one-illusion* blocks, in each of which the test phase consisted of one object size embedded in one of the two illusory configurations in 20 consecutive trials (such that both configurations were grasped in this type of block), two *two-sizes, one-illusion* blocks with 20+20 trials of two different objects (40 and 45 mm, presented in an intermixed, pseudo-randomised order) embedded in the same illusion configuration, and one *one-size, two-illusions* block of 20+20 trials of the same object embedded in two different illusion configurations (again intermixed, pseudo-randomised). The trial order was pseudo-randomised by using cycles containing all combinations of size*illusion configuration in random, not predetermined order. The order of blocks was determined by using a row-balanced combination of two 5*5 Latin squares. This resulted in a total of 280 trials for each participant, with a Müller-Lyer illusion presented in 160 of them. Prior to the first block, participants were given an undetermined small number (<10) of practice blocks with vision of their hand, before the mirror's occluding back panel was in place. They were then given an unspecified number of further practice trials, until they felt comfortable executing the task. During practice, no illusion display was presented, and trials were not recorded.

2.3 Data processing and analysis

Kinematic data were analysed in R (R Core Team, 2015). For each grasp, we extracted the MGA, response time, movement time, and time to MGA from the raw, unfiltered frame-wise data. The start of the grasping movement was determined through a position criterion relative to the starting position (thumb and index finger > 30 mm away from starting position), while the end of the grasp was determined through an aperture criterion (within 10 mm of the object size, aperture velocity < 5 mm/sec for at least 60 msec).

Our main dependent variable was the MGA. Thus, we submitted all MGAs to a 2*3*2 repeated-measures ANOVA with factors *object size, illusion angle* (which had two levels representing the illusion configurations, and one level representing the lack of illusory fins), and *experiment subset* (first vs. second half of trials). Greenhouse-Geisser epsilons and corrected p-values (Greenhouse & Geisser, 1959) were computed for factors with more than two levels. This was followed by an analysis of the overall illusion effect, as well as the illusion effect over bins of trials, which we computed as units comprising of two trials per configuration, the reason being that this was the smallest unit for which an illusion effect could be computed while still allowing for single invalid trials. Binning the illusion effect in such a way allowed us to then run a linear regression of illusion effect over bins to investigate a possible decrease in illusion effect over time (as was done in Franz et al., 2001).

The illusion effect is typically computed as the difference of MGA in mm between grasps toward objects within the incremental (fins outwards) illusion and grasps toward objects within the decremental (fins inwards) illusion. To be able to compute a within-subject illusion effect over time, that is, for individual bins in which the number of trials in each size and configuration may not be balanced (e.g., due to single missing trials), we used an additional layer in which we calculated the mean MGA difference between grasps toward objects embedded in the incremental illusion and non-illusory objects of the same size, and the mean difference between grasps toward non-illusory objects and objects within the decremental illusion. The sum of these differences was used as the overall illusion effect. A slope-corrected illusion effect was calculated in % of object size divided by the responsiveness of grasping (see Bruno & Franz, 2009; Franz, Scharnowski, & Gegenfurtner, 2005), primarily to enable an easier comparison to other studies, as slope-correction is not necessary for comparisons within one measure. We calculated the corrected effect with the formula

$$(1)\ i_{corr} = \frac{i_{raw}}{s} * \frac{100}{l}$$

and the standard error by using a Taylor-approximation (see Franz, Hesse, & Kollath, 2009),

$$(2)\ SE_{icorr} = \frac{i_{raw}}{s} * \sqrt{\frac{\sigma_s^2}{s^2} + \frac{\sigma_i^2}{i_{raw}^2} - \frac{2\sigma_{is}}{i_{raw} * s}} * \frac{100}{l}$$

with i_{corr} = the corrected illusion effect, i_{raw} = the raw illusion effect, s = the mean slope (responsiveness) of the MGA, l = the mean object size (length), σ_s = the SEM of the slope, σ_i = the SEM of the raw illusion effect, σ_{is} = the between-subject covariance between the illusion effect and the MGA-slope.

2.4 The error correction model

The process of visuomotor adaptation studies is commonly formalized with linear state-space models (Cheng & Sabes, 2006; Thoroughman & Shadmehr, 2000). Briefly, the point of this type of model is to update an internal state on a trial-by-trial basis according to experienced movement errors (or two states, see e.g. Taylor, Wojaczynski, & Ivry, 2011). The main goal of our study was modelling the effect of such a mechanism on grasping a visual illusion and propose a mechanism for the sometimes-observed motor learning in such designs. As was done by Cesanek et al. (2016), we modelled grasp planning with a linear function that maps perceived sizes onto MGAs, where the intercept parameter is a dynamic internal state (x_n) and the fixed slope (α; Säfström & Edin, 2005) is estimated for each individual by linear regression of object sizes onto MGAs observed during baseline. Thus, the planned MGA would be calculated as follows:

$$(3)\ MGA_{planned} = x_n + \alpha * (l + \omega),$$

with x_n = an internal state, α = the slope of the response function, l = object length, and ω representing a visual perturbation induced by the illusion. States were updated based on two processes, (1) error correction and (2) decay of the previous state. This was modelled using the following equation:

$$(4a)\ x_{n+1} = a * x_n - bC * e_n,$$

where a = a retention parameter indicating the stability of the state over consecutive trials or the rate the state 'decayed' toward the initial state, x_n = the state for the n-th trial, bC = an error correction parameter representing the amount of learning from error in previous trials, and e_n = the error on the n-th trial, defined simply as the difference between the planned MGA (modelled from the current state and the visual object size) and the mean MGA. Note that such an error would be observed upon touching the object, not visually during the movement, since participants had no vision of their hands. To model the situation where two opposite illusion configurations were presented in interleaved order, we expanded the model to include two separate states x_n^D and

x_n^I for the decremental and incremental illusion configuration, respectively, as well as separate error terms e_n^D and e_n^I and an error generalization parameter bG:

$$\text{(4b)} \begin{bmatrix} x_{n+1}^D \\ x_{n+1}^I \end{bmatrix} = \begin{bmatrix} a & 0 \\ 0 & a \end{bmatrix} * \begin{bmatrix} x_n^D \\ x_n^I \end{bmatrix} - \begin{bmatrix} bC & bG \\ bG & bC \end{bmatrix} * \begin{bmatrix} e_n^D \\ e_n^I \end{bmatrix}.$$

On each trial, only one illusion configuration was presented; thus, one state may be considered the ‘active’ state and one the ‘inactive’ state, and one of the error terms (corresponding to the ‘inactive’ state) is set to 0. Multiplying the matrices shows that each state is updated to the previous state for this particular configuration multiplied by a, minus either (a) the error term multiplied by bC (for the ‘active’ state) or (b) the error term multiplied by bG (for the ‘inactive’ state). Thus, the error generalization parameter indicates the degree to which error in one type of illusion configuration updates the state used for grasp planning toward the other configuration.

To be able to compute a planned MGA for the first trial, we also included the initial state as a parameter in the model. This results in four free parameters in the case of one-illusion blocks (retention parameter a, error correction parameter bC, initial state x_0, and perturbation ω), and six in the two-illusion block (an additional error correction generalization parameter bG, and two perturbation parameters ω^D and ω^I instead of one). These parameters were bounded to be within realistic ranges – for details, as well as an example of how a state and the predicted MGA would be calculated, see the appendix.

These parameters can be interpreted in a very straightforward way: The perturbation parameters represent the (not measurable) underlying perceptual illusion effect in mm, the error correction parameters quantify what proportion of movement planning error is corrected on a given trial, and the decay parameter indicates the proportion of adaptation that is carried over to the next trial (thus slowing down learning when it is < 1, as well as giving a natural asymptote for the measured illusion effect over time). They also serve to express most of our predictions. Firstly, we expected positive error correction parameters, both for bC (indicating participants correcting previous over- or underestimation of object size) and bG (indicating interference between illusion conditions, such that an observed error in the opposite illusion on the previous trial might still affect the next trial, only in the wrong direction). We also expected ω^D and ω^I to be negative and positive, respectively, if the model were to faithfully represent the illusion. We had no predictions for parameters a and x_0.

3. Results

In the full group of 47 participants, 768 trials (5.8%) were missing due to the specified maximum of two times that an invalid trial could be repeated; however, the vast majority of these (741 trials) were among the participants excluded due to missing data, see section 2.1 and footnote 1. A total of 47 trials (0.4% of all trials) were removed from analyses as outliers due to not having a valid MGA (i.e., the measured MGA being smaller than the object presented), and another 68 trials (0.6%) for being greater than 4 SD removed the participant’s mean MGA. This is a slightly more liberal criterion than typically used – this is because much of our analyses relied on individual trials instead of cell means, so we wanted to eliminate only obviously extreme outliers. These exclusion criteria left us with 40 participants with data that allowed us to compute illusion effects over time (section 2.1, footnote 1). Of the data from these participants, 89 trials total (0.8%) were either not completed (27 trials) or had to be removed from analysis (62 trials). On average, participants reached the target after 1019 msec, with the MGA occurring on average after 58% of the movement time. Using MGA as the dependent variable, we then conducted a 2*3*2 repeated-measures ANOVA on the remaining data, which revealed a main effect of factors *object size* ($F(1, 39) = 42.41$, $p < .001$), *fin angle* ($F(2, 78) = 24.29$, $\varepsilon_{gg} = .89$, $p_{gg} < .001$), as well as *subset* ($F(1, 39) = 8.40$, $p = .006$). A single interaction was statistically significant, *fin angle * subset* ($F(2, 78) = 4.40$, $\varepsilon_{gg} = .99$, $p_{gg} = .016$), showing that the effect of the illusion configuration differed by subset, very much in line with what we expected to find and replicating findings by Whitwell et al. (2016) in the Ponzo illusion. All other interactions were non-significant (all $p > .4$). Overall, the results from our ANOVA show a clear influence of the Müller-Lyer illusion on the MGA, with a possibility of a decrease that we explored in further analyses.

3.1 The illusion effect over time

To quantify the Müller-Lyer-illusion effect on the MGA, we first calculated an overall illusion effect as the mean difference between grasps toward targets with an incremental illusory context and size-matched targets with a decremental illusory context (see section 2.3). We then applied a correction for the responsiveness of participants' MGA (mean slope: 0.69, with an S.E.M. of 0.13) and calculated the effect in % of object size (see equations 1 and 2, as well as Bruno & Franz, 2009; Franz et al., 2009, 2005). This gave us an overall illusion effect of 13.3 per cent larger MGAs for incremental vs. decremental illusion trials, 95% CI [6.0, 20.6], which was statistically significant ($t(39) = 3.68$, $p < .001$) and falls well within the range of the effects found in Bruno and Franz's (2009) review.

Next, we calculated the illusion effect separately for each block, and for each bin within each block. This was done to enable us to look at the magnitude of the illusion effect over time and test for a possible decrease, as well as potential differences between blocks with regards to the dynamics of the illusion. We chose to collapse trials over bins instead of on a trial-by-trial basis to ensure that (a) data from all participants were included for each data point, allowing us to compute within-participant effects, and (b) object sizes were always balanced. The results are plotted in Figure 2 (mean MGAs by trial can be seen in Figure 3). In statistical terms, we did see a decrease in the illusion effect: Over all data, linear regression[2] showed that the illusion effect decreased by an average of 0.33 percentage points per bin, which was a statistically significant decrease ($t(39) = -3.042$, $p = .004$). Splitting up the data by blocks revealed a negative slope in two of three blocks, that is, single-illusion, single size (mean slope: -0.45 percentage points/bin; paired t-test for difference from 0: $t(39) = -2.587$, $p = .013$), single-illusion, dual-size (-0.36 %/bin; $t(39) = -2.162$, $p = .037$), but not dual-illusion (-0.13 %/bin; $t(39) = -1.455$, $p = .154$). However, these slopes were not significantly different from one another (all $p > .12$). The illusion

[2] Note that we would expect a nonlinear decrease of the illusion effect, which is also predicted by the error-correction model. Linear regression slopes are included to give the reader a sense of the magnitude of adaptation on an easily interpretable scale, and to allow some rough inference statistics along with the more detailed modelling.

effects by subset can also be seen in Figure 2 (right), illustrating the interaction found in our ANOVA.

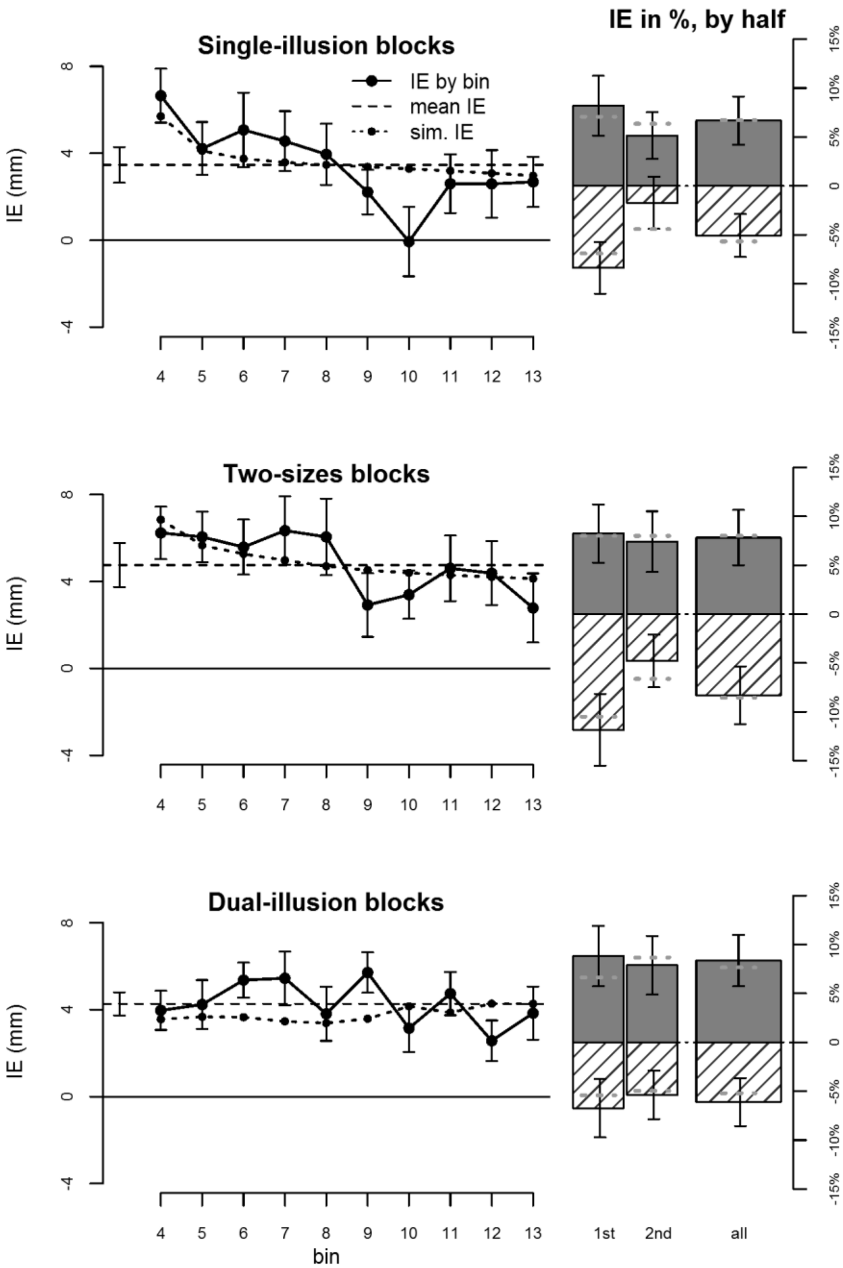

Figure 2: Measured and simulated illusion effects over time. **Left**: Mean illusion effect (black) for each bin, in mm, by block type. Dashed line indicates overall mean illusion effect. Dotted line indicates illusion effect computed from simulated MGAs as described in section 2.4. Bin count starts at 4 due to the baseline phase. Error bars indicate within-subject S.E.M. for the pooled differences between bins (Franz & Loftus, 2012; Loftus & Masson, 1994), allowing inferences about the dynamics of the illusion effect. Bars on the dashed line indicate S.E.M. of the illusion effect. **Right**: Bars show observed mean corrected illusion effect in per cent ± within-subject S.E.M. (dark grey: incremental, hatched: decremental), for first and second half of each block, and overall. Grey dotted horizontal lines indicate simulated corrected illusion effects, split up in the same way.

3.2 The dynamic illusion effect as error correction

We have shown evidence that the illusion effect decreases over trials. To strengthen this assertion and offer an explanation why this may be the case, we fit a model of error correction to the data (see section 2.4, specifically equations 3 and 4). In essence, this model fit a number of parameters (four, if there was only one illusory context: An initial state, a retention parameter, an error correction parameter, and the perturbation; six if there were two illusory contexts, with an additional error generalization parameter and a second perturbation), bounded to reasonable limits, on a per-subject and per-block basis. The model minimized the root mean squared error (RMSE) of the model-simulated MGAs vs. observed MGAs. The mean MGAs predicted by the model, as well as the mean observed MGAs, are plotted in Figure 3.

Table 1: Mean parameter values and goodness of fit.

block	a	bC	bG	ω^I	ω^D	x_0	RMSE
Dual-illusion	0.91 [0.89, 0.92]	0.14 [0.03, 0.27]	0.12 [0.03, 0.25]	1.14 mm [-0.44, 2.75]	-1.46 mm [-3.10, 0.24]	47.88mm [37.52, 58.47]	6.06 mm
Single, increment	0.88 [0.75, 0.90]	0.19 [0.09, 0.31]	NA	2.31 mm [0.86, 3.42]	NA	49.08mm [38.88, 60.27]	5.75 mm
Single, decrement	0.84 [0.75, 0.90]	0.29 [0.17, 0.43]	NA	NA	-1.34 mm [-2.63, -0.23]	45.09mm [35.18, 55.24]	5.44 mm
Dual-size, increment	0.89 [0.81, 0.93]	0.17 [0.07, 0.31]	NA	2.41 mm [1.15, 3.57]	NA	50.49mm [39.99, 61.91]	6.13 mm
Dual-size, decrement	0.92 [0.84, 0.95]	0.19 [0.01, 0.32]	NA	NA	-0.98 mm [-1.97, 0.04]	45.19mm [35.73, 54.98]	6.63 mm

Note: Parameter values given with 95% confidence intervals obtained via BCa bootstrap. Parameters indicate retention (a), error correction (bC), error generalization, or interference (bG), the visual perturbations corresponding to the illusion (ω^D and ω^I), and the initial state (x_0). RMSE indicates root mean-squared error of simulated vs. actual MGA.

Mean parameter values and RMSEs for each of the five blocks can be seen in Table 1, along with confidence intervals obtained by estimating S.E.M.s via BCa-bootstrap (Efron & Tibshirani, 1993). In contrast to the linear regression reported above (section 3.1), which gives a sense of the overall reduction in the illusion effect over time, the state-space model quantifies the proportion of experienced error that is corrected following each movement. It can be seen that the retention parameter a was close to 1 in all blocks – indicating that the state did not change much apart from the correction for an error signal. For our question of whether error correction could explain dynamic illusion effects in grasping, the error correction parameter bC was central, which indicated the degree to which participants responded to being 'off the mark' and corrected their error. This parameter was significantly different from 0 in all blocks, indicating that under the model, error correction may explain the decreasing effect. As expected, ω^D parameters were in all cases smaller than 0 in the expected direction, while ω^I parameters were larger than 0. These parameters rather closely matched the measured uncorrected illusion effects of -1.6 mm and 2.33 mm. In fact, as the illusion effect estimated by the model fit was 4.2 mm, compared to the observed overall raw illusion effect of 3.93 mm.

As it has sometimes been reported that the illusion effect (Kopiske et al., 2016) or the decrease in the illusion effect (Whitwell et al., 2016) can be asymmetric, we compared the two parts of the illusion effect (that is, the mean MGA difference between each illusion configuration and a neutral stimulus) with a paired t-test, finding no significant difference ($t(39) = 0.908$, $p = .370$). The error correction parameters were also very similar (0.18 and 0.24 respectively), although visual inspection reveals a possibly slightly larger effect in the incremental illusion (Figure 3). However, the only conclusion our data allow with regards to such an asymmetry effect is that it cannot be very large, if it exists.

For our hypothesis that the number of illusion configurations would impact the degree of adaptation, and the notion that potentially the number of objects per block would also matter for this, two differences were critical. The former factor was easily tested by looking at the error generalization parameter (which indicated interference between opposing illusion configurations): Parameter bG was significantly higher than 0, indicating that error generalization took place. In other words, learning of the decremental effect of inward-pointing fins interfered with learning of the incremental effect of outward-pointing fins. For the latter factor (number of objects), we had no modelling parameter, so we looked at the difference in error-correction bC

parameters between blocks using one object size and blocks using two objects sizes. No such differences were statistically significant, and indeed no such trend was visible, either, so there was no support for this hypothesis.

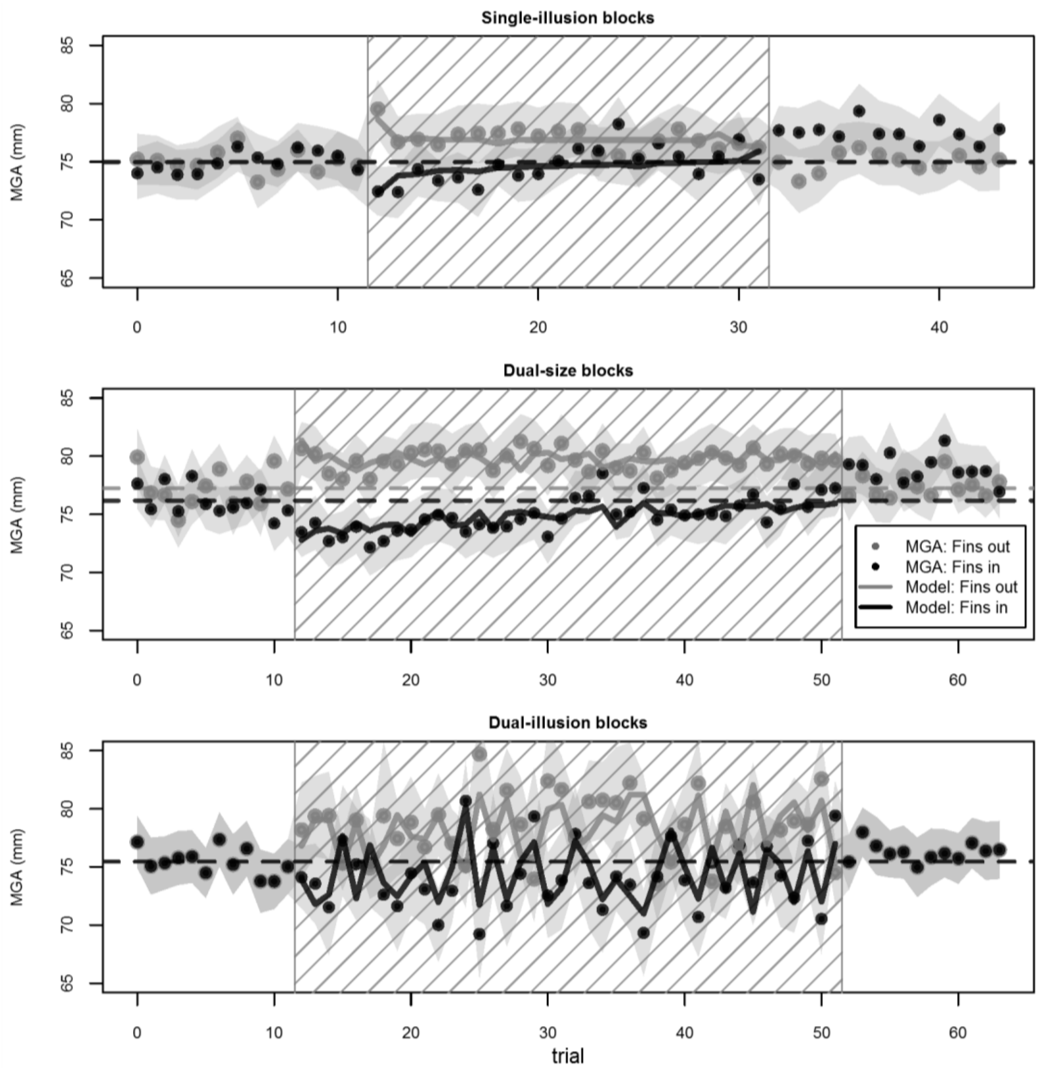

Figure 3: MGA by trial. Mean MGA in mm, plotted over trials, split up by block type. Dots indicate actual mean MGA (grey for incremental, black for decremental illusion trials or blocks) ± S.E.M. indicated by grey area. Hatched area indicates illusion phase. Lines indicate predictions of the model with best-fitting parameters, see Table 1. Dashed lines indicate mean responses in baseline and washout phases. Responses to different sizes are collapsed in all graphs. The anti-phases in the dual-illusion blocks are a direct consequence of the cycle-constrained pseudo-randomisation.

Beyond the model, we also tested for other characteristics of sensorimotor adaptation and error correction. A key prediction of modelling grasping illusion as error correction would be that following a trial where the target appeared larger than it was, the MGA should be noticeably smaller to correct for the error presumably made in the previous trial. Indeed, this was the case, as the mean MGA[3] in trials following incremental illusion trials was on average 0.77 mm smaller than following decremental illusion trials, a statistically significant difference ($t(39) = 3.372$, $p = .002$) – see also Figure 4. The figure also shows that as we expected, the MGA in a given trial was not inversely affected by the object size in the previous trial. In fact, larger objects in a previous trial did tend to produce a larger MGA in the current trial ($t(39) = 2.787$, $p = .008$).

Central to the idea of adapting is also an aftereffect: That is, immediately after the perturbation is removed we should see an effect opposite to the perturbation. Indeed, we found that comparing MGAs in the first trial of washout phases following single-illusion blocks revealed a 3.39 mm larger MGA in washout trials following the decremental illusion than the incremental illusion ($t(39) = 2.306$, $p = .026$). A comparison of the first trials of each washout phase to the baseline

[3] We accounted for the fact that incremental illusion trials were disproportionally likely to be followed by decremental illusion trials, and vice-versa, by first computing means for each previous*current illusion type combination, and then averaging these.

phase, as is commonly done in the adaptation literature (e.g., Fernández-Ruiz & Díaz, 1999), could not be performed confound-freely due to the fact that MGA was generally larger in the second half of each block.

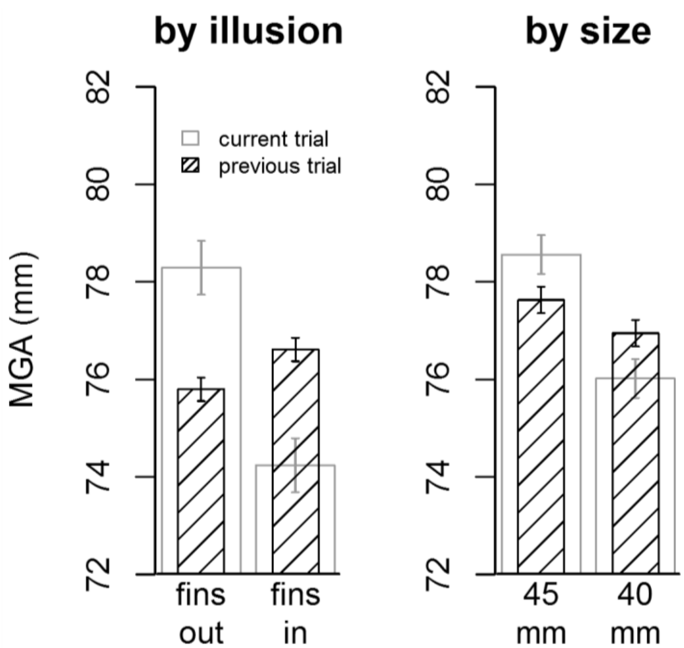

Figure 4: Effects of previous trials on MGA. MGA by size and illusion context of the previous (black, hatched) and current trial (grey, unfilled). As expected, objects within the incremental illusion were grasped larger than objects within the decremental illusion (left, grey), and larger objects were grasped larger than smaller objects (right, grey). Crucially, as black hatched bars show, trials following an incremental illusion were grasped smaller (left, black hatched), indicating an opposite effect consistent with error correction. This opposite effect was not present for object sizes (right, black hatched). Again, a second layer was included to aggregate data as described in footnote 3. Figure displays data from phase two of two-illusion blocks and two-size blocks only. Error bars indicate the within-subject SEM of the difference between the two conditions (see Franz & Loftus, 2012; Loftus & Masson, 1994).

4. Discussion

Our results show two things quite clearly: One, that grasping is susceptible to the Müller-Lyer illusion, consistent with the review by Bruno and Franz (2009). The illusion effect on grasping is strong and robust, and unquestionably present in all blocks of the experiment. Two, the effect decreases with repeated trials, replicating recent findings in the Ponzo illusion (Cesanek et al., 2016; Whitwell et al., 2016), and extending findings that had found adaptation to the Müller-Lyer illusion in saccades (Bruno, Knox, & de Grave, 2010; Knox, 2010).

4.1 Modelling the illusion as error correction

Our findings can be explained well in terms of an error-correction model of sensorimotor adaptation with an error signal felt upon touching the object, which may be a step towards understanding *how* adaptation to visual illusions might work. We also found a decrease for the effect of both the incremental and the decremental illusion display. The results are less clear regarding the question whether the number of different stimulus conditions in a block affects how quickly participants can adapt to the illusion. Our results clearly show interference between multiple illusion configurations, which was one of our main predictions, although inference statistics on linear regression slopes did not find a pronounced difference in the degree of adaptation across single-illusion and dual-illusion blocks, and there was virtually no sign of an effect of presenting multiple sizes in a single block. More general principles of error correction also fit well with our data. Primarily, there is the interdependence between trials: Incremental and decremental illusion trials led to an opposite adjustment of MGA in the trials that followed, and also to aftereffects in the washout phase. Finally, the model also does not assume that the different sizes introduced in some blocks would make any difference at all. We nevertheless tested for some effect on the adaptation rate, as laid out in section 1.2 – but the data were in agreement with the model, as neither comparison of regression slopes, nor of error correction parameters showed any difference between single-size and dual-size blocks.

To fit our model, one difficulty in treating visual illusions as size perturbations is that we can only have a vague idea of the degree to which the illusion perturbs participants' perception. Indeed, we know that participants vary in how much they are affected by visual illusions (see, e.g., Coren & Porac, 1987). This makes it necessary to fit the illusion on an individual basis, as we did. In fact, each parameter was fit not only for each participant, but also for each block separately, allowing another plausibility check. Parameters were consistent across blocks in that parameter values for the illusion effects and correction parameters were in the expected direction and of roughly similar size (see Table 1), strengthening our confidence that the model provided a sensible fit. Our model needs few parameters to achieve a reasonably good fit: An initial state, retention parameter and error correction parameter are enough to capture the change in the illusion effect in single-illusion blocks. In dual-illusion blocks, we also fit a generalization parameter, although this was not strictly required; this parameter was our way of testing whether the additional configuration could explain differences in illusion-effect dynamics in previous studies. Indeed, we found a positive generalization parameter, indicating that interference in adaptation was caused by error correction without specificity for illusory context. A comparison by Akaike's information criterion slightly preferred the simpler model (ΔAIC of -2.2 for the simpler model over the more complex one that included a bG parameter; see Akaike, 1974; Burnham & Anderson, 2004), indicating that the improvement in fit was quite small. Given that the more complex model is the only one that can explain both the adaptation aftereffect in single-illusion blocks and the previous-trial effect in dual-illusion blocks, we still prefer it.

As illustrated in Figure 3, the model captured quite nicely the general trend of the nonlinear decreasing illusion effect in single-illusion blocks (similar to what has been found in previous studies Buckingham & Goodale, 2010; Cesanek et al., 2016; Whitwell et al., 2016), as well as the oscillating and more slowly decreasing illusion effect in the dual-illusion blocks. It is also important to note that by including a baseline phase before introducing the illusion, as is typically done in research on visuomotor adaptation, participants got to execute a number of 'real', non-practice grasps before we started measuring the illusion, thus providing another safeguard to ensure that they were not still learning to execute the task at this stage.

4.2 Our study and the perception-action debate

Our results speak strongly against a perception-action model in which visual processing for skilled actions is uniformly unaffected by contextual illusions (Aglioti et al., 1995; Milner & Goodale, 1995, 2006). However, Whitwell and colleagues (2016) recently argued that a similar result with the Ponzo illusion does not necessarily conflict with this model. Finding a decrease in the Ponzo illusion effect on grasping, they proposed that this could be due to the grasp control system initially treating the illusion context as an obstacle, but learning that it is not a real obstacle. Our results cast doubt on this interpretation: in the decremental context, avoidance of the inward-pointing fins could not possibly produce a smaller grasp. Whitwell and colleagues (2016) go on to suggest that it might be precisely the fact that actions adapt to illusions that distinguishes vision-for-action from vision-for-perception. They propose that a decreasing illusion effect is the consequence of a sensorimotor system that uses visual and haptic information to "refine the programming of visually guided grasping" (Whitwell et al., 2016, p. 1163). Whitwell et al. (2016) leave open some specifics of how this adaptation of vision-for-action might operate, but it is clear that it would be a (a) property of dorsal processing that (b) relies on some dynamic updating. Here, we have shown that if we assume a dynamic visuomotor mapping (a standard assumption in the visuomotor adaptation literature, see e.g., Ghahramani et al., 1996; Shadmehr et al., 2010; Taylor et al., 2011), the additional assumption of two separate size representations is not needed. Thus, our model should be considered more parsimonious than a modified perception-action model.

Indeed, we show that in our data, grasping a visual illusion behaves just like grasping a visual perturbation induced by a prism or a mirror, where participants absolutely cannot see the true size of an object: Using a common baseline-perturbation-washout design we find the classic pattern of (i) an initially large but decreasing effect of the perturbation (i.e., the illusion), (ii) a washout effect once it is removed (see Figure 3), as well as (iii) a previous-trial effect when opposite illusions are interleaved. For this larger point it is not important whether we measure grip aperture or grip force (Hesse et al., 2016), or whether a model considers object size or contact points (Smeets & Brenner, 1999, 2006): We argue that treating visual illusions like regular visual perturbations is a fruitful and parsimonious approach to understanding their effects on action.

We summarise the main implications of our results and our model for the grasping-illusion debate as follows: Is possible to explain differences in illusion effects between perception and grasping without assuming two different visual representations. In fact, our model explicitly assumes grasping responses to be based on consistent, stable perceptual illusion effects. In abstract terms, this is a specific case of the more general principle that aggregation in repeated-measures designs can be problematic if the effect to be measured is not constant over time. We also show that interleaving multiple illusion configurations can hinder, but perhaps not always fully prevent learning of the illusion, even in a setup where different illusions were not spatially separated. On the other hand, using two object sizes rather than only one had no measurable effect on motor learning. Thus, interference in the absence of spatial separation may account for the fact that illusion effect did not decrease in studies like Franz et al. (2001) but did in others (this study; Cesanek et al., 2016; Whitwell et al., 2016), while the different number of objects seems to be irrelevant. Future studies may be useful to verify the importance of spatial separation, as well as investigating differences between the smaller, but still existent decrease despite interference in our study and the virtually constant illusion effect in others. In our view, sensorimotor adaptation is a promising framework for such investigations that can, as demonstrated, predict and explain illusion effects in grasping while making plausible and relatively few assumptions.

5. Conclusion

We show that a standard model of sensorimotor adaptation through error correction qualitatively predicts the dynamics of an illusion effect on grasping by treating the visual size illusion as a simple perturbation of size perception. Such a model could provide a plausible, theory-driven approach for integrating the literature.

Acknowledgements

The authors would like to thank Siyu Chen, Bethany Hung, Carly Kaplan, and Aubryn Samaroo for their help collecting the data.

References

Aglioti, S., DeSouza, J. F. X., & Goodale, M. A. (1995). Size-contrast illusions deceive the eye but not the hand. *Current Biology*, *5*(6), 679–685.

Akaike, H. (1974). A new look at the statistical model identification. In *IEEE Transactions on Automatic Control AC-19* (pp. 716–723).

Bruno, N., & Franz, V. H. (2009). When is grasping affected by the Müller-Lyer illusion? A quantitative review. *Neuropsychologia*, *47*, 1421–1433.

Bruno, N., Knox, P. C., & de Grave, D. D. J. (2010). A metanalysis of the effect of the Müller-Lyer illusion on saccadic eye movements: No general support for a dissociation of perception and oculomotor action. *Vision Research*, *50*(24), 2671–2682. http://doi.org/10.1016/j.visres.2010.09.016

Buckingham, G., & Goodale, M. A. (2010). The influence of competing perceptual and motor priors in the context of the size-weight illusion. *Experimental Brain Research*, *205*(2), 283–288. http://doi.org/10.1007/s00221-010-2353-9

Burnham, K. P., & Anderson, R. P. (2004). Multimodel Inference: Understanding AIC and BIC in Model Selection. *Sociological Methods & Research*, *33*(2), 261–304. http://doi.org/10.1177/0049124104268644

Cesanek, E., Campagnoli, C., & Domini, F. (2016). One-shot correction of sensory prediction errors produces illusion-resistant grasping without multiple object representations. *Journal of Vision*, *16*, 20. http://doi.org/10.1167/16.12.20

Cheng, S., & Sabes, P. N. (2006). Modeling sensorimotor learning with linear dynamical systems. *Neuronal Computation*, *18*(4), 760–793. http://doi.org/10.1162/089976606775774651.Modeling

Coren, S., & Porac, C. (1987). Individual differences in visual-geometric illusions: Predictions from measures of spatial cognitive abilities. *Perception & Psychophysics*, *41*(3), 211–219. http://doi.org/10.3758/BF03208220

Efron, B., & Tibshirani, R. J. (1993). *An introduction to the bootstrap*. New York: Chapman & Hall.

Fernández-Ruiz, J., & Díaz, R. (1999). Prism adaptation and aftereffect: Specifying the properties of a procedural memory system. *Learning & Memory*, *6*(1), 47–53. http://doi.org/10.1101/lm.6.1.47

Flanagan, J. R., & Beltzner, M. A. (2000). Independence of perceptual and sensorimotor predictions in the size-weight illusion. *Nature Neuroscience*, *3*(7), 737–741.

http://doi.org/10.1038/76701

Franz, V. H., Fahle, M., Bülthoff, H. H., & Gegenfurtner, K. R. (2001). Effects of visual illusions on grasping. *Journal of Experimental Psychology: Human Perception and Performance*, *27*(5), 1124–1144.

Franz, V. H., Gegenfurtner, K. R., Bülthoff, H. H., & Fahle, M. (2000). Grasping visual illusions: No evidence for a dissociation between perception and action. *Psychological Science*, *11*(1), 20–25.

Franz, V. H., Hesse, C., & Kollath, S. (2009). Visual illusions, delayed grasping, and memory: No shift from dorsal to ventral control. *Neuropsychologia*, *47*, 1518–1531.

Franz, V. H., & Loftus, G. R. (2012). Standard errors and confidence intervals in within-subjects designs: Generalizing Loftus and Masson (1994) and avoiding the biases of alternative accounts. *Psychonomic Bulletin & Review*, *19*, 395–404.

Franz, V. H., Scharnowski, F., & Gegenfurtner, K. R. (2005). Illusion effects on grasping are temporally constant not dynamic. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(6), 1359–1378.

Ghahramani, Z., Wolpert, D. M., & Jordan, M. I. (1996). Generalization to local remappings of the visuomotor coordinate transformation. *The Journal of Neuroscience*, *16*(21), 7085–7096.

Glover, S., & Dixon, P. (2001). Motor adaptation to an optical illusion. *Experimental Brain Research*, *137*(2), 254–258. http://doi.org/10.1007/s002210000651

Goodale, M. A. (2014). How (and why) the visual control of action differs from visual perception. *Proceedings of the Royal Society B*, *281*, 20140337. http://doi.org/http://dx.doi.org/10.1098/rspb.2014.0337

Goodale, M. A., Meenan, J. P., Bülthoff, H. H., Nicolle, D. A., Murphy, K. J., & Racicot, C. I. (1994). Separate neural pathways for the visual analysis of object shape in perception and prehension. *Current Biology*, *4*(7), 604–610.

Goodale, M. A., Milner, A. D., Jakobson, L. S., & Carey, D. P. (1991). A neurological dissociation between perceiving objects and grasping them. *Nature*, *349*, 154–156.

Greenhouse, S. W., & Geisser, S. (1959). On methods in the analysis of profile data. *Psychometrika*, *24*, 95–112.

Hesse, C., Miller, L., & Buckingham, G. (2016). Visual information about object size and object position are retained differently in the visual brain: Evidence from grasping studies. *Neuropsychologia*, *91*, 531–543. http://doi.org/10.1016/j.neuropsychologia.2016.09.016

Jackson, S. R., & Shaw, A. (2000). The Ponzo illusion affects grip-force but not grip-aperture scaling during prehension movements. *Journal of Experimental Psychology: Human*

Perception and Performance, *26*(1).

Jeannerod, M. (1984). The timing of natural prehension movements. *Journal of Motor Behavior*, *16*(3), 235–254.

Jeannerod, M. (1986). The formation of finger grip during prehension. A cortically mediated visuomotor pattern. *Behavioural Brain Research*, *19*, 99–116.

Keefe, B. D., & Watt, S. J. (2009). The role of binocular vision in grasping: A small stimulus-set distorts results. *Experimental Brain Research*, *194*(3), 435–444. http://doi.org/10.1007/s00221-009-1718-4

Knox, P. C. (2010). The reduction of the effect of the Müller-Lyer illusion on saccade amplitude by classic adaptation. *I-Perception*, *1*(2), 95–102. http://doi.org/10.1068/i0395

Kopiske, K. K., Bruno, N., Hesse, C., Schenk, T., & Franz, V. H. (2016). The functional subdivision of the visual brain: Is there a real illusion effect on action? A multi-lab replication study. *Cortex*, *79*, 130–152. http://doi.org/10.1016/j.cortex.2016.03.020

Krakauer, J. W., Ghez, C., & Ghilardi, M. F. (2005). Adaptation to Visuomotor Transformations: Consolidation, Interference, and Forgetting. *Journal of Neuroscience*, *25*(2), 473–478. http://doi.org/10.1523/JNEUROSCI.4218-04.2005

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review*, *1*, 476–490.

Milner, A. D., & Goodale, M. A. (1995). *The visual brain in action* (1st ed.). Oxford: Oxford University Press Inc.

Milner, A. D., & Goodale, M. A. (2006). *The visual brain in action* (2nd ed.). Oxford: Oxford University Press Inc.

R Core Team, T. (2015). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from https://www.r-project.org

Säfström, D., & Edin, B. B. (2004). Task requirements influence sensory integration during grasping in humans. *Learning & Memory*, *11*, 356–363.

Säfström, D., & Edin, B. B. (2005). Short-term plasticity of the visuomotor map during grasping movements in humans. *Learning & Memory*, *12*(1), 67–74. http://doi.org/10.1101/lm.83005

Schenk, T., Franz, V. H., & Bruno, N. (2011). Vision-for-perception and vision-for-action: Which model is compatible with the available psychophysical and neuropsychological data? *Vision Research*, *51*, 812–818.

Shadmehr, R., Smith, M. A., & Krakauer, J. W. (2010). Error correction, sensory prediction, and adaptation in motor control. *Annual Review of Neuroscience*, *33*, 89–108.

Smeets, J. B. J., & Brenner, E. (1999). A new view on grasping. *Motor Control*, *3*(3), 237–271.

Smeets, J. B. J., & Brenner, E. (2006). 10 Years of Illusions. *Journal of Experimental Psychology: Human Perception and Performance*, *32*(6), 1501–1504.

Soechting, J. F., & Flanders, M. (1989). Sensorimotor representations for pointing to targets in three-dimensional space. *Journal of Neurophysiology*, *62*(2), 582–594.

Taylor, J. A., Wojaczynski, G. J., & Ivry, R. B. (2011). Trial-by-trial analysis of intermanual transfer during visuomotor adaptation. *Journal of Neurophysiology*, *106*(6), 3157–3172. http://doi.org/10.1152/jn.01008.2010

Thoroughman, K. A., & Shadmehr, R. (2000). Learning of action through adaptive combination of motor primitives. *Nature*, *407*, 742–747.

von Helmholtz, H. (1867). *Handbuch der physiologischen Optik*. Leipzig, Germany: Leopold Voss.

Westwood, D. A., & Goodale, M. A. (2011). Converging evidence for diverging pathways: Neuropsychology and psychophysics tell the same story. *Vision Research*, *51*, 804–811.

Whitwell, R. L., Buckingham, G., Enns, J. T., Chouinard, P. A., & Goodale, M. A. (2016). Rapid decrement in the effects of the Ponzo display dissociates action and perception. *Psychonomic Bulletin & Review*, *23*(4), 1157–1163. http://doi.org/10.3758/s13423-015-0975-4

Woodworth, R. S. (1899). The accuracy of voluntary movement. *Psychological Review-Monograph Supplements*, *3*(3), 1–114.

Woolley, D. G., Tresilian, J. R., Carson, R. G., & Riek, S. (2007). Dual adaptation to two opposing visuomotor rotations when each is associated with different regions of workspace. *Experimental Brain Research*, *179*(2), 155–165. http://doi.org/10.1007/s00221-006-0778-y

Ypma, J. (2014). nloptr: R interface to NLopt. Retrieved from https://cran.r-project.org/package=nloptr

Appendix: Details of the error correction model

The model was implemented using the nloptr package for R (Ypma, 2014), which used the cobyla algorithm to minimize the RMSE between the actual and simulated MGA. Simulated MGA was calculated as simMGA = slope * visualSize + x_n, with visualSize = physicalSize + ω. This was done on a per-block, per-participant basis. The optimization ran until iterations changed each parameter by less than 0.01%, or after a maximum of 10^16 iterations, using the starting values and bounds given in Table A.1.

Table A.1: Parameters used in the model.

Parameter	Description	Starting value	Bounds
$\boldsymbol{x_0}$	Initial state	50	[0, 120]
$\boldsymbol{a}$	Retention	0.95	[0, 1]
$\boldsymbol{bC}$	Error correction	0	[-1, 1]
$\boldsymbol{bG}$	Error generalization	0	[-1, 1]
$\boldsymbol{\omega^I}$	Perturbation, fins out	Mean raw IE, fins out	[-10, 10]
$\boldsymbol{\omega^D}$	Perturbation, fins in	Mean raw IE, fins in	[-10, 10]

Note: Initial state $\boldsymbol{x_0}$ bounded generously between the smallest and the largest intercept observed (rounded to the nearest multiple of 10). $\boldsymbol{a}$ bounded between 0 (never departing from previous state) and 1 (instant decay toward $\boldsymbol{x_0}$). bC and bG between -1 and 1, representing 100% error correction and 100% error carry-over, respectively. ω bounded generously between realistic illusion effect values.

So, what do these values mean? Figure A.1 displays some examples of how different parameter values might impact the simulated MGA. Let us also consider the case of a dual-illusion block with a 40-mm object, a mean observed response function of y = 48.78 + 0.69x, and the mean parameter values to see what the model looks like in action. In this case, we might start with a state of x_0 = 47.88 and a decremental illusion trial. Filling in the parameter values, we get a simulated MGA of

$$simMGA_1 = 0.66 * 38.54 + 47.88 = 73.32,$$

which we compare to the predicted MGA for an unperturbed object to get the error signal:

$$e_1^D = 71.77 – (0.66 * 40 + 48.78) = -1.86.$$

Note that the error signal is much more strongly negative than the perturbation due to the x_0 parameter. Indeed, it is also possible to get a negative error signal despite a positive perturbation – a fact that may be counterintuitive, but is not problematic or worrying since it is very plausible that participants might be subject to certain biases at the start of the illusion phase, but serves to illustrate that the perturbation and the error signal are distinct. We use this error signal (and ignore e_1^I, as it is 0) to compute the states for the next trial:

$$x_1^D = a * (x_0 – x_0) + x_0 – bC * e_1^D = 0.91 * 0 + 47.88 – 0.14 * -1.86 = 48.14,$$

and

$$x_1^I = a * (x_0 – x_0) + x_0 – bG * e_1^D = 0.91 * 0 + 47.88 – 0.12 * -1.86 = 48.10.$$

So, a negative error signal (a too small grasp) is corrected for, leading to a larger state for both configurations – but more strongly so for the decremental one, since $bC > bG$. We also see in these calculations how the retention parameter a is applied not to the state, but the difference between the state and x_0 – this difference is going to be non-zero in the following trials. This ensures that the state decays toward the initial state, not toward 0. Continuing in such a way will produce a sequence of states. These along with object sizes and illusion types can be used to simulate a sequence of MGAs, which is exactly what our model did.

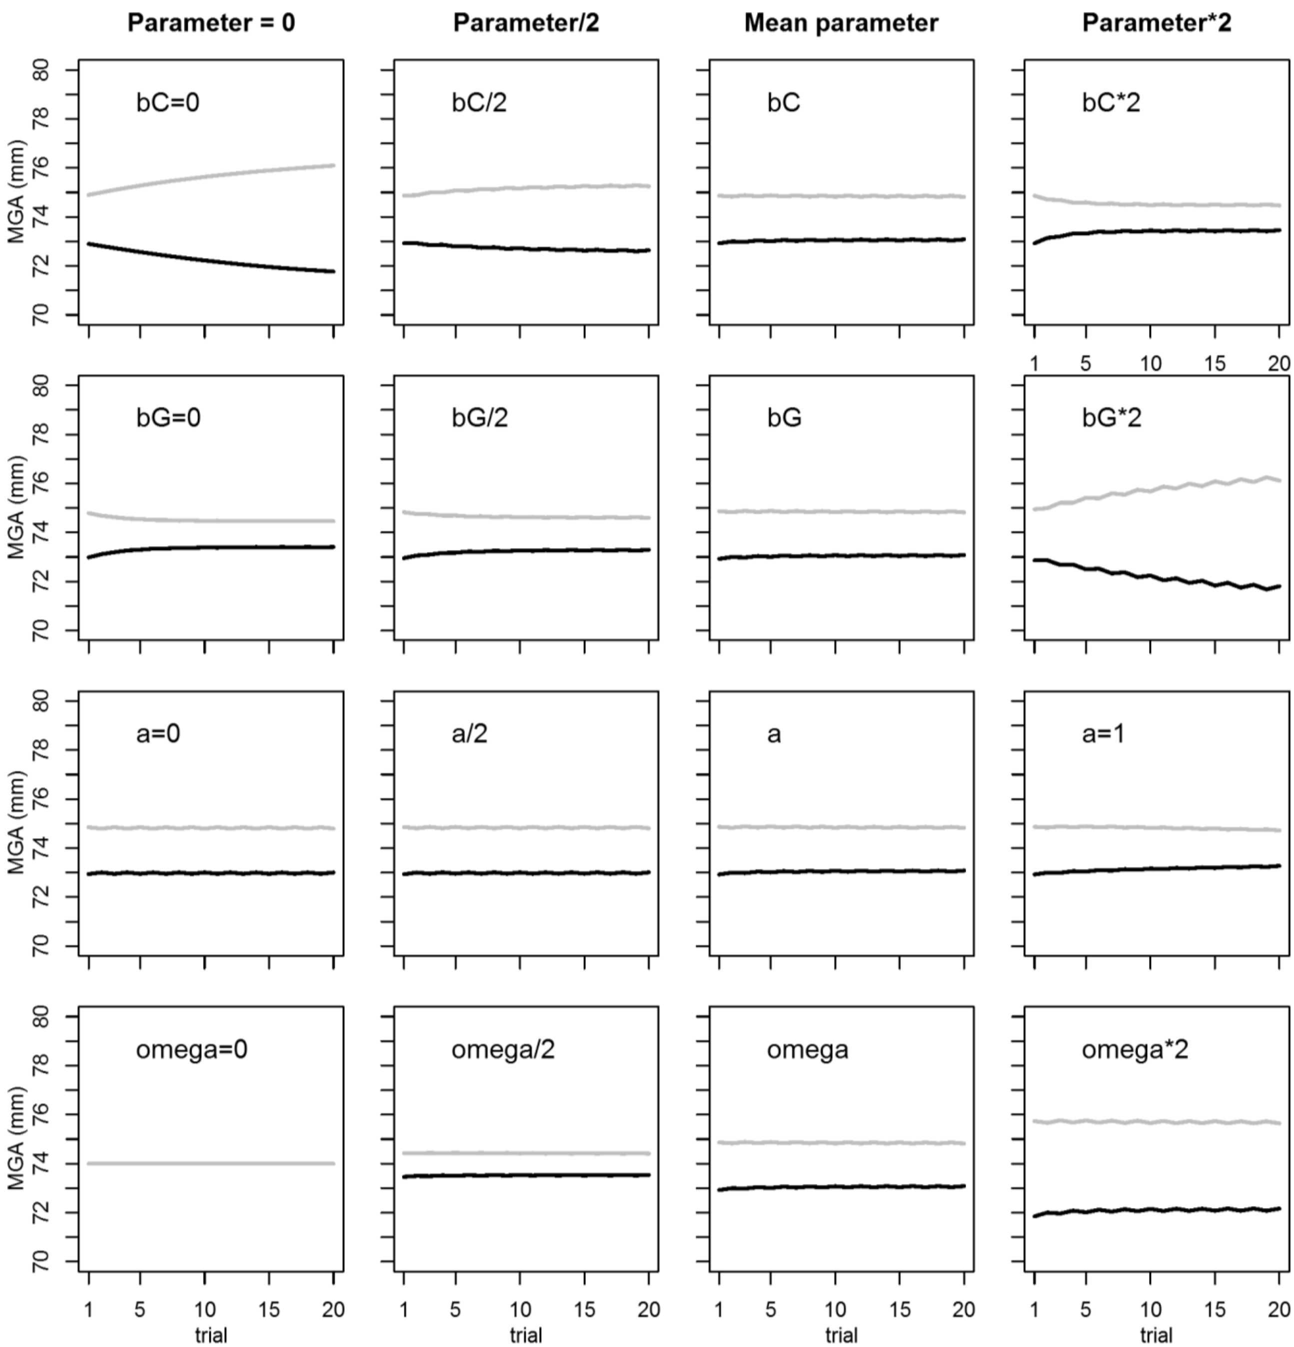

Figure A.1: Simulated MGAs with sample parameter values. Simulated MGAs plotted over trials. Grey = fins out, black = fins in. First row: Varying parameter *bC*. Second row: Varying *bG*. Third row: Varying *a*. Fourth row: Varying ω. Parameters are set to 0 (left column), half the mean value obtained in our data (second from the left), the actual mean value obtained in our data (second from the right), and twice the obtained mean value or their upper bound (right column). Each data point based on the model applied to 10,000 simulated sequences. The anti-phase is caused by the pseudo-randomization procedure and can also be seen as in figure 3.