Drawing Accuracy, Quality and Expertise

by

Linda Carson

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Psychology

Waterloo, Ontario, Canada, 2012

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Drawing from a still-life is a complex visuomotor task. Nevertheless, experts depict three-dimensional subjects convincingly with two-dimensional images. Drawing research has previously been limited by its general dependence on qualitative assessment of drawings by human critics and on retrospective self-report of expertise by drawers. Accuracy measures cannot hope to encompass all the properties of "goodness" in a drawing but this thesis will show that they are consistent with the expertise of the drawers and with the quality ratings of human critics, they are robust enough to support analysis of ecologically valid drawing tasks from complex three-dimensional stimuli, and they are sensitive enough to study global and local properties of drawings.

Drawing expertise may depend to some extent on more accurate internal models of 3D space. To explore this possibility we had adults with a range of drawing experience draw a still life. We measured the angles at intersecting edges in the drawings to calculate each person's mean percentage magnitude error across angles in the still life. This gave a continuous objective measure of drawing accuracy which correlated well with years of art experience. Participants also made perceptual judgements of still lifes, both from direct observation and from an imagined side view. A conventional mental rotation task failed to differentiate drawing expertise. However, those who drew angles more accurately were also significantly better judges of slant, i.e., the pitch of edges in the still life. Those with the most drawing experience were significantly better judges of spatial extent, i.e., which landmarks were leftmost, rightmost, nearest, farthest etc.. The ability to visualize in three dimensions the orientation and relationships of components of a still life is related to drawing accuracy and expertise.

In our second study, we set out to extend our understanding of drawing accuracy and to develop measures that would support more complex research questions about both drawing and visual perception. We developed and applied novel objective *geometric* measures of accuracy to analyze a perspective drawing task. We measured the deformation of shapes in drawings relative to the ground truth of a reference photograph and validated these measures by showing that they discriminate appropriately between experts and novices. On all measures—orientation, proportionality, scale and position of shapes—experts out-performed novices. However, error is not uniform across the image. Participants were better at capturing the proportions and positions of objects (the "positive space") than of the spaces between those objects (the "negative space") and worse at orienting those objects than shapes in the negative space, but scale error did not differ significantly between positive and negative space. We have demonstrated that objective geometric measures of drawing accuracy are consistent with expertise and that they can be applied to new levels

of analysis, not merely to support the conventional wisdom of art educators but to develop new, evidence-based means of training this fundamental skill.

Most or all prior research into drawing was based on human ratings of drawing quality, but we cannot take for granted that the "goodness" of a drawing is related to its accuracy. In order to determine whether our objective measures of accuracy are consistent with drawing quality, we invited more than one hundred participants to grade the quality of all of the drawings we had collected and measured. We showed participants photographs of the still lifes on which the drawings were based and asked them to grade the quality of each drawing on a scale from 1 ("Poor") to 10 ("Excellent"). People's quality ratings were consistent with one another. People without drawing experience rated drawings slightly more highly than the drawing experts did, but the ratings of both groups correlated well. As we predicted, the more drawing experience the artist had, the more highly rated the drawing was, and the more accurate the drawing was, the more highly rated it was. Furthermore, scaling error (but not proportionality, orientation or position) also predicted drawing quality. In perspective drawing, accuracy—as measured by angle error or polygon error—is related to drawing quality.

If drawing practice strengthens an artist's model of 3D space, we would expect the three-dimensionality of drawings to be disrupted by damage to the dorsal stream or the connection between the dorsal and ventral streams. A former illustrator and animator, DM, who had suffered a right hemisphere stroke and presented with spatial neglect, performed modified versions of the angle judgement, spatial judgement and indirect drawing tasks of our second study. Despite his previous experience, he showed weaknesses in his mental model of 3D space, weaknesses that were not evident in his drawings before the stroke.

Taken together, the thesis has developed and validated two objective measures of drawing accuracy that both capture expert/novice differences well and provide superior measures when contrasted with self-reported expertise. The performance of a single patient with neglect highlights the potential involvement of the dorsal stream in drawing. The novel quantitative measures developed here allow for testable hypotheses concerning the cognitive and neural mechanisms that support the complex skill of drawing to be objectively measured.

Acknowledgements

First and foremost, I gratefully acknowledge the wisdom and resilience of my supervisor, James Danckert, who welcomed me to his lab with a prediction that I embraced as a mandate: “At least you won’t be boring.” Thank you, James, for leading by example. I acknowledge the helpful scrutiny of my entire committee, who improved this work the good old-fashioned way, by making me revise and defend it. I acknowledge the continued mentorship of Fran Allard, whose supervision of my M.Sc. laid the foundation for this work and whose classroom teaching set a standard I strive for every day. I acknowledge the generous, practical advice and commiseration of my dear friends and sometime collaborators: Matthew Millard, who set the pace, and Parthipan Siva, who crossed the finish line with me. I acknowledge the support of labmates present and past, who modeled excellence and collegiality for me every day. I acknowledge the good-natured co-operation of my participants in the face of unusually difficult drawing tasks.

Anne Lamott wrote, “The reason life works at all is that not everyone in your tribe is nuts on the same day.” Get through it however you can. I got through it with #madwriting, #FridayNightScience, Dr. V., Gmail chat, Skype, Dropbox, Beanstalk and the friends of “Shut Up and Write.”

I have been lucky to have the help of some hard-working UW staff and research assistants—Bill Eickmeier and Marg Ingleton; Inara Aliev, Evan Borman, Emma Cotten and Nadine Quehl—some very kind friends—Betty Bax, Nina Brickman, Anna Foat, Jim Gardner, Kate Hoye, Cate Huston, Darlene McGeer, Prabhakar Ragde and Sherrie Steinberg—and my unexpected parents—Bill and Veronica Carson.

My thinking is usually improved by exposure to my students. Thank you to the students of Fine Arts, Independent Studies and Knowledge Integration at the University of Waterloo. Most of all, I acknowledge the formative influence of my students at the big black pig studio, whose trials and triumphs regularly inspired me to wonder “How *do* we learn to draw?”

Dedication

This is dedicated to all of my students, but especially to everyone who learned to draw, or to draw better, at the big black pig studio.

Table of Contents

List of Tables

List of Figures

Chapter 1

Introduction

Drawing is a complex visuomotor skill that has attracted research interest in recent years (Kozbelt & Seeley, 2007; Mitchell, Ropar, Ackroyd, & Rajendran, 2005; Tchalenko & Chris Miall, 2009). That research has been limited, however, by its general dependence on qualitative assessment of drawings by human critics (Cohen, 2005). These measures have reasonable utility (i.e., ratings are validated by tests of inter-rater reliability; Chan, Chan, & Chau, 2009; Kozbelt, 2001) and undoubtedly capture a many-dimensioned understanding of "goodness" in drawings. However, they are global (rating the drawing as a whole) and coarse-grained (rating drawing performance on a few-point scale), which drastically limits the experimental questions that can be asked and answered.

This thesis is motivated by a lifelong interest in how we teach, learn and practise drawing, and a desire to use drawing to extend our knowledge of human visual perception. To do that, we need a yardstick.

The accuracy of a drawing should be more practical to measure—albeit more demanding—than its quality, though accuracy measures will not capture all of the aesthetic virtues of a perspective drawing (nor any of the virtues of an abstract drawing). Some recent work in computer graphics, for example, measures drawing accuracy to support research interests such as where an artist would draw a line to depict a particular image in order to develop improved sketching interfaces for graphical software. Schmidt and colleagues (2009) asked participants to draw projections of curves in 3D scenes then assessed their accuracy by visual inspection, looking for specific biases in perspective. Cole and colleagues (2012) compared artists' line drawings to a tonal reference image. The property of interest was where artists chose to put particular lines to represent gradual changes in form and an accuracy-like measure was used to gauge how well the participants' drawings agreed with one another. For key points on the image, they measured how often the artists' lines were within 1mm of each other.

Such piece-by-piece measures of mark-making accuracy tell only part of the story of how we draw. Researchers into human behaviour and perception have also investigated the accuracy of line drawing. Cohen and Bennett (1997) compared novice's freehand drawings from photographs with tracings of these photographs. When accuracy was rated by human critics, the traced photographs were superior. This suggests that even when people are demonstrably able—when tracing—to make all the marks necessary for a drawing, they often fail, drawing freehand, to put those marks together with good judgment.

Tchalenko (2007) demonstrated that experts and novices were equally capable of drawing isolated line segments with reasonable fidelity. Interestingly, a more recent study (Tchalenko, 2009) demonstrated that when required to draw those same line segments now embedded within an entire drawing, accuracy decreased for the same novices. In other words, copying a simple line was much easier than copying the same line embedded within a non-trivial pictorial array. Imagine, then, how much more difficult is the task of the professional artist, who not only makes those lines but creates the entire pictorial array from observation of the three-dimensional subject. Mark-making accuracy is, perhaps, analogous to how success in writing the letters of the alphabet well is necessary but not sufficient for successfully writing words and paragraphs.

This thesis is built on three drawing experiments and a related case study. The study in Chapter 2 develops an angle-based measure of drawing accuracy, validates it against the experience levels of the participants, and provides evidence of a top-down influence of drawing context on local accuracy. Angle error gives a basis for studying accuracy locally within the image. Comparing drawing accuracy with performance on several novel perceptual tasks also demonstrates that drawing skill is linked to a better mental model of three-dimensional space. The study in Chapter 3 develops a geometric approach to drawing accuracy, trading the relative simplicity of angle error calculations for a more versatile, powerful polygon-based measure. This measure is also consistent with the drawing experience of the participants. Polygon error gives a basis for studying drawing regions of particular research interest, such as objects. It separates drawing error into four independent properties—orientation, proportionality, scale and position—and establishes that experts are more accurate than novices on all of them. Participants in general–regardless of experience–draw more accurate proportionality and position in positive space than in negative space, but more accurate orientation in the negative space. In the study in Chapter 4, human critics rate the overall *quality* of the same drawings measured in Chapters 2 and 3, confirming that both angle-based and polygon-based measures of accuracy are consistent with subjective quality ratings. This is important to position the findings in the context of the drawing literature. Finally, Chapter 5 describes the performance of an artist with neglect on several of these experimental tasks. Despite his previous drawing experience, his judgements of angles and spatial extent were largely indistinguishable from novices, and his drawing showed both mild neglect of the left side of the image and signs of impairment to his 3D mental model of space.

1.1 Expertise

In an expertise approach (Ericsson & Charness, 1994), researchers compare what experts and novices do in the course of completing domain-specific and domain-general tasks, in order to isolate performance differences that spotlight which practices are genuinely central to expertise. Expert performance differences have been found in domains as varied as chess (Gobet & Simon, 1996), sight-reading music (Lehmann & Ericsson, 1996), volleyball (Allard & Starkes, 1980) and solving physics problems (Chi, Glaser, & Rees, 1982; VanLehn & Van, 2009).

The first challenge to the researcher is how to discriminate experts from novices. Researchers have identified experts based on competitive success (Halpern & Wai, 2007), professional credentials (Wei & Luo, 2010), post-secondary education (Sanjram & Khan, 2011) or performance on an objectively-measured behavioural task (Gauthier, Skudlarski, Gore, & Anderson, 2000). Conventionally, in the absence of a performance-based measure of expertise, participants are designated as experts if they have ten or more years of training and experience. This standard has often been sufficient to detect meaningful, informative differences between experts and novices (Ericsson, Krampe, & Tesch-Römer, 1993).

1.2 Perspective drawing

Much (but not all) drawing is representational—the drawing is an image that represents a concrete subject such as a person, a landscape or a vase of flowers. When the artist attempts to accurately represent not just the subject but the geometric specifics of how the subject would really look from a particular viewpoint, that is perspective drawing. To draw in perspective, an artist inspects (or invents) a three-dimensional subject and draws on the page her best approximation of how the subject projects onto the two-dimensional "picture plane" (a fronto-parallel plane between the view and the subject). The ability to draw in perspective has been a desirable skill for artists since at least the Renaissance

(Bermingham, 2000; Goldstein, 1996; Kubovy, 1986). Today, even artists working in other media and styles usually spend years as students mastering perspective drawing as a fundamental part of art education (Nicolaides, 1941).

The research in this thesis is a rare example of studying drawings made from observation of a three-dimensional subject, not from a photograph (Cohen, 2005), an artist's drawing (Tchalenko, 2009), a diagram (Mitchell et al., 2005), or geometric instructions (Tchalenko, 2007). Only a handful of studies have investigated drawings from 3D (Howard & Allison, 2011; Kozbelt, 2001; Miall & Tchalenko, 2001), and arguably none of them dealt with complex subjects of non-trivial depth. Howard and Allison (2011) asked participants to complete a line drawing of a cube. Kozbelt (2001) had people draw from observation of a pair of scissors. The most complex subject described in the literature was probably a portrait drawn from life (Miall & Tchalenko, 2001) but the properties of research interest in that instance were the eye and hand movements of the artist. Even in that case, the subject was unoccluded and occupied a relatively shallow pictorial space. Perspective drawing is more demanding when the subject has many elements, occlusion and non-trivial depth, as in the experimental tasks here.

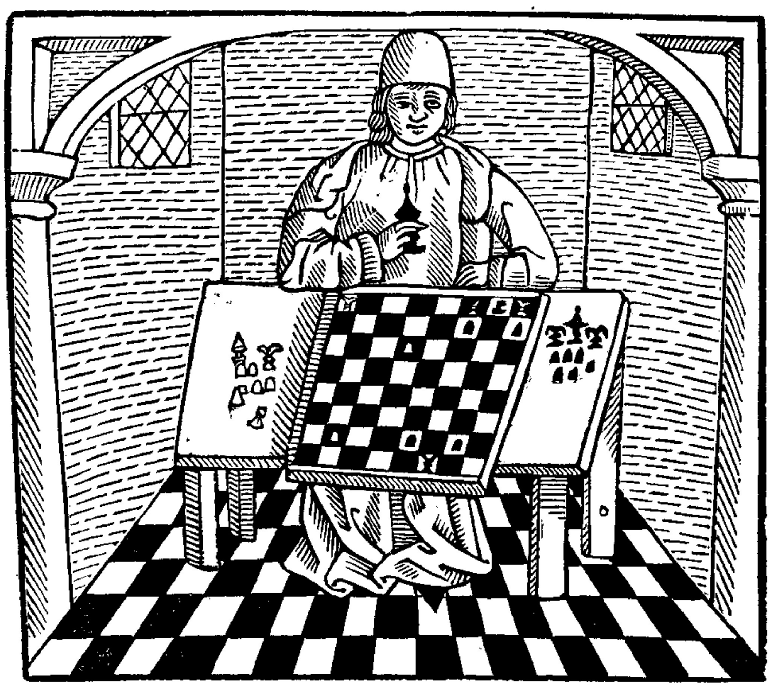

Figure 1.1: The table and chessboard in this medieval illustration (artist unknown, from Caxton (1474)) show divergent perspective.

There are some biases in perspective drawing well-known to artists and researchers alike. For example, neither children nor adults foreshorten the receding edges of a cube sufficiently (Nicholls & Kennedy, 1995). Howard and Allison (2011) recently showed that when people were asked to complete the top face of a perspective drawing of a cube (from direct observation of a 3D cube or by copying a 2D drawing of a cube) they had a strong tendency to draw the receding edges as divergent when, in fact, those edges should appear to converge. This corresponds to the tendency, commonly seen in drawings, to draw the tops of objects as if from a bird's eye view, a tendency common to modern students (Lockard, 1982; Roth, 1992) and medieval artists (Figure 1.1). Analysis of local drawing accuracy should be able to detect and quantify such perspective errors.

Figure 1.2: This Renaissance drawing device (Dürer & Strauss, 1525/1977) was used to depict the subject in strictly accurate perspective. The artist, standing right, studies the subject, seated left, by tracing his image onto the vertical "window" between them. The window is a physical analogue of the "picture plane." To maintain a fixed viewpoint, the artist sights along the (adjustable) vertical rod. The orientation of each edge the artist draws on the window is analogous to what was investigated in this thesis as *projected angle*.

In our studies, we asked participants to draw a three-dimensional still life (and to make perceptual judgements of it) from a specific viewpoint, and coached them on which landmarks to align in order to find and maintain that viewpoint. We did not physically constrain them, such as with a chin rest. We strongly prefer this approach in order to keep the methods as similar to domain-specific drawing tasks as possible. Many expertise researchers have demonstrated the importance of using ecologically valid tasks, that is, tasks which capture as much of natural performance conditions as possible (e.g., Mann, Abernethy, & Farrow, 2010). During perspective drawing, artists must regularly solve the challenge of maintaining—and restoring—a consistent viewpoint (Figure 1.2). We expect that the ability to do so will vary with the expertise of the participants; inaccuracy caused by any failure to maintain a consistent viewpoint is part of what we should be measuring, not controlling, in this research.

1.3 Measuring drawing accuracy

In general, drawing research relies almost exclusively on the ratings of human critics (Cohen, 2005). Even in computer graphics, drawing researchers have used human critics to assess accuracy (e.g., Schmidt et al., 2009). Where accuracy is measured objectively, it has been measured in specific regions of the drawing rather than overall. For example, Cole and colleagues (2012) measured whether artists' lines were within 1mm of one another at key points.

Ideally, an objective measure of drawing accuracy would be practicable, would improve the reliability and statistical power of analyses, would eliminate the quest for expert raters (or the defense of the use of novice raters), and would afford analysis of properties of both regions of a drawing (local) and the drawing as a whole (global).

For a measure of drawing accuracy to be credible, it should be consistent with the level of experience of the drawer. Nevertheless, we should expect to see some differences between a specific performance-based measure of skill and a general self-report of domain experience. Self-report and retrospection are problematic, and the ten-year threshold for expertise is a rough cut. Drawing experience includes more skills than accuracy, art experience includes more than drawing, and drawing includes more than perspective. Imagine an art student humbly under-reporting her years of practice, and a veteran abstract painter who has spent most of his years on colour theory rather than perspective drawing. Despite the difference between their years of art education and professional experience, there is reason to wonder which of them is truly a drawing expert. A performance-based measure should group each of them with other participants of similar levels of demonstrated skill rather than by seniority.

Chapter 2

Angle accuracy and a mental model of three-dimensional space

2.1 Introduction

What makes expert artists experts? Drawing “in perspective” from a still life is a complex visuomotor task. Nevertheless, experts convincingly depict three-dimensional subjects with two-dimensional images.

In previous work (Carson & Allard, 2008, in press), we developed the use of angle accuracy to investigate perspective drawing. Participants drew from direct observation of a three-dimensional still life, seen from a controlled viewpoint, and we compared selected angles in those drawings to the same angles as measured on a ground truth photograph of the still life taken from the same viewpoint. We chose this approach for three reasons. First, drawing from a still life (rather than from a photograph or another drawing) is a complex, ecologically valid task that should inform us about real-world drawing behaviour and challenge even expert participants. Second, the accuracy of an angle is independent of its size, so we could allow participants to draw freely at any scale that was comfortable for them. Finally, angle judgement was a promising candidate for a relevant perceptual skill when looking for an expertise effect in perspective drawing. In the expertise approach, it is economical to study performance on similar but contrasting tasks, one that is arguably domain-specific and another that is arguably domain-general. The characteristic pattern of results in other domains of expertise is that domain-specific tasks discriminate between experts and novices, while similar domain-general tasks do not. The accuracy of a drawing is clearly a domain-specific performance measure, and we believed—correctly, as it turned out—that we could measure a drawing's accuracy by measuring angles within it. Verbal judgement of the accuracy of planar angles is probably not a domain-specific task. To investigate the logical connection between angle judgement and drawing, we also asked participants to make angle judgements of edges in the still life itself. That is, we asked them to judge the slant of edges in the still life and we asked them to judge the angle at which an edge should be drawn. We speculated that drawing experts might have a discernible advantage over novices on the latter task where, arguably, angle judgement and drawing meet. Our findings did not support that hypothesis, but—unexpectedly—did suggest experts might be better than novices at judging slant in the still life. In the present study, we set out to replicate and extend that work.

Drawing expertise may depend on more accurate mental models of 3D space. To test this possibility, we asked adults to draw from a still life and to make perceptual judgements of still lifes, both from direct observation and as they *imagined* the still life would look from a side view. We studied their performance on the perceptual tasks in the context of both their drawing *experience*—a self-report measure of expertise in drawing—and their *accuracy* in drawing the still life in perspective—a performance-based measure of drawing skill.

A high priority was making the tasks ecologically valid in order to maximize our likelihood of capturing evidence of the domain-specific perceptual demands of drawing. To that end, we made the drawing task as authentic an experience as possible, and made its still life subject complex enough to challenge even our most experienced participants. We embedded the new perceptual tasks in a similarly complex, drawing-like context. Ecological design considerations included making the drawing subject three dimensional rather than asking participants to copy a photograph or another drawing, providing a binocular view, and coaching them on how to maintain the correct viewpoint rather than constraining them physically.

In this study, we compared a conventional ten-year expert/novice criterion to a novel, objective measure of drawing accuracy (based on angle drawing error) Such a measure offers three advantages. It is based on performance rather than reputation and self-report, it distinguishes between more finely-grained levels of expertise, and it analyzes drawing accuracy locally within the drawing, as well as providing a global measure of overall performance. We validated this measure of drawing accuracy against the standard ten-year expert/novice measure and used it to investigate our main research question: *Do people who draw more accurately also judge 3D space more accurately?*

Measuring a drawing's accuracy by measuring its angles can be done simply with a protractor and has the benefits of being objective and independent of drawing size. For example, Howard and Allison (2011) used participants' angle error when completing the edges of the top and side faces to investigate the use of divergent perspective in cube drawing. In 1943, Cain measured very simple drawings of irregular hexagons copied from planar images and reported the magnitude error of the angles of the figures. His angle error results correlated with instructor ratings of participants' performance in a drawing class, and a later test showed that senior art students outperformed the first-year students. McManus and colleagues (2010) recently used a task based on Cain's work to show that drawing ability is related to the ability to copy simple angles and proportions.

In this study, we measured angle accuracy locally throughout the drawing, took a mean across the entire image as a measure of overall drawing accuracy, and asked participants to make angle judgements of several kinds. The study of angle judgement is a rich and complicated area of scientific literature. Researchers as early as Jastrow (1892) reported as conventional wisdom that people regularly underestimate acute angles and overestimate obtuse angles . . . then produced experimental evidence to the contrary. Pratt (1926) and Beery (1968), in turn, countered Jastrow based on different experimental tasks. As the evidence accumulated, it became clear that angle judgement varied systematically with angle size but in patterns more fine-grained than "acute/obtuse" (Chen & Levi, 1996). Some angle sizes may be special (e.g., ninety degrees, Gray & Regan, 1996; Ferrante, Gerbino, & Rock, 1995) though that too has been contested (Snippe & Koenderink, 1994). Very small and very large angles are subject to floor and ceiling effects (MacRae & Loh, 1981). Angle judgement varies with the design of the task (Wenderoth, OConnor, & Johnson, 1986), response modality (Maclean & Stacey, 1971), and even the age of the participant (Simon & Ward, 1979). There are systematic patterns in people's accuracy when judging and reproducing angles, but research has yet to isolate all of the moderating factors.

The present study differs from the majority of the angle judgement literature in one especially important property: We are not studying angle judgements or drawings made in isolation. We are asking participants to extract angle information from a complex, layered and deep three-dimensional subject. We are measuring the angle information embedded in drawings, which are complex arrays of many neighbouring angles constructed in context and in sequence.

We report all our angle errors—drawing and judgement—as the magnitude of the difference between the reported angle and the correct value, scaled to the percentage of the correct value. For example, if the participant drew an angle at 27° which was, in the ground truth photograph of the still life, 30°, the percentage magnitude angle drawing error would be $\left|\frac{30-27}{30}\right| = 10\%$. We use magnitude rather than signed error in all our angle measures. This is consistent with the work of Cain (1943) and McManus and colleagues (2010), and ensures that the errors are unaffected by the choice of which of a pair of complementary angles we measure in a drawing. Furthermore, while the direction of the error is often of interest in other psychological research, our previous work (Carson & Allard, 2008, in press) showed that experts made significantly smaller errors than novices on some angle tasks and found no noteworthy expert/novice differences on the *direction* of those errors. We scale all our angle errors to a percentage of the correct values. An error of two or three degrees is probably of little concern with respect to an angle of 120° but very important indeed for an angle of just 10°.

In this study, we wanted to probe for evidence that people who draw well in perspective also have a more accurate mental model of three-dimensional space. This research question was triggered by previous work (Carson & Allard, 2008, in press) but also by anecdotal evidence from the drawing studio, where some veteran art instructors appear to be able to visualise the still life or the life model from students' vantage points around the room. This could be the expression, in the studio, of people's general ability to mentally rotate three-dimensional objects (Shepard & Metzler, 1971), but it may also be an indication that experienced artists are better at visualising the objects in space than are novices. Perez-Fabello and Campos (2007) recently showed that art training was associated with the ability to visualise and draw a room, based on a written description, as it would appear to someone standing at the door and also to mentally transform the scene, drawing the room as it would appear through the window. Fifth-year art students were better at both tasks than first-year art students, which also suggests this is a skill improved by art training. While Rock, Wheeler, and Tudor (1989) reported that participants were very poor at drawing a twisted wire shape from an imagined side view, they also mentioned that two experienced drawers produced "quite accurate drawings" when they tried the task informally.

In this study, we looked for domain-general skill in mental rotation with a conventional pencil-and-paper mental rotation task (Peters et al., 1995). We looked for domain-specific skills in inspecting the subject—here, the still life—and understanding its spatial properties by asking participants to make perceptual judgements about a still life (slant, projected angles, occlusion and spatial extent). We also combined those themes with a quixotic investigation of domain-specific mental rotation, asking participants to make perceptual judgements of one still life and to draw another *as each still life would look from an imagined viewpoint.*

2.2 Method

2.2.1 Participants

We recruited 35 adult participants (nine men, 26 women; age $M = 27.09$, $SD = 13.5$) to make perspective drawings and perceptual judgements from complex three-dimensional still lifes.

Most of the participants who were recruited from the university student population earned course credit for their participation. We also recruited participants from the local art community in order to find drawing experts. In the expertise approach, it is conventional—in the absence of a performance-based measure—to designate participants as expert if they have ten or more years of training and experience in the domain of interest (Ericsson et al., 1993). Based on a survey of the participants' educational and professional history in drawing, nine were identified as experts (four men, five women). All of the expert participants were at least 20 years of age, and none of the novice participants was over 29.

2.2.2 Apparatus

Complex three-dimensional still lifes A still life is a tabletop arrangement of objects used as a subject for drawing from direct observation. While the popularity of the still life as a subject of commercial or critical interest in art varies, the still life is a perennial staple of drawing instruction and practice (Edwards, 1979). For this study, the central tasks were to make drawings and perceptual judgements from complex three-dimensional still lifes. For those tasks, we prepared two still lifes of comparable complexity, each thoroughly but unobtrusively marked so we could reconstruct it reliably for each participant. We named each still life for our own records ("Bravo," and "Charlie," Figure 2.4) but did not use those names with the participants.

Each still life was presented on a table in a well-lit room while the participant sat roughly two metres away. A researcher sat to the right of the participant, facing the side of the still life, also roughly two metres away from the table (Figure 2.1). Participants used a 2B pencil (and eraser) on an $8\frac{1}{2}$ by 11 inch piece of commonplace 20-pound white paper to draw, supported by a clipboard.

Each still life included a small visual landmark we used to coach participants on how to locate and maintain the desired viewpoint consistently (similar to Figure 2.2). When the participant's view was correctly aligned, the landmark was occluded, to avoid making the landmark more salient than its neighbourhood.

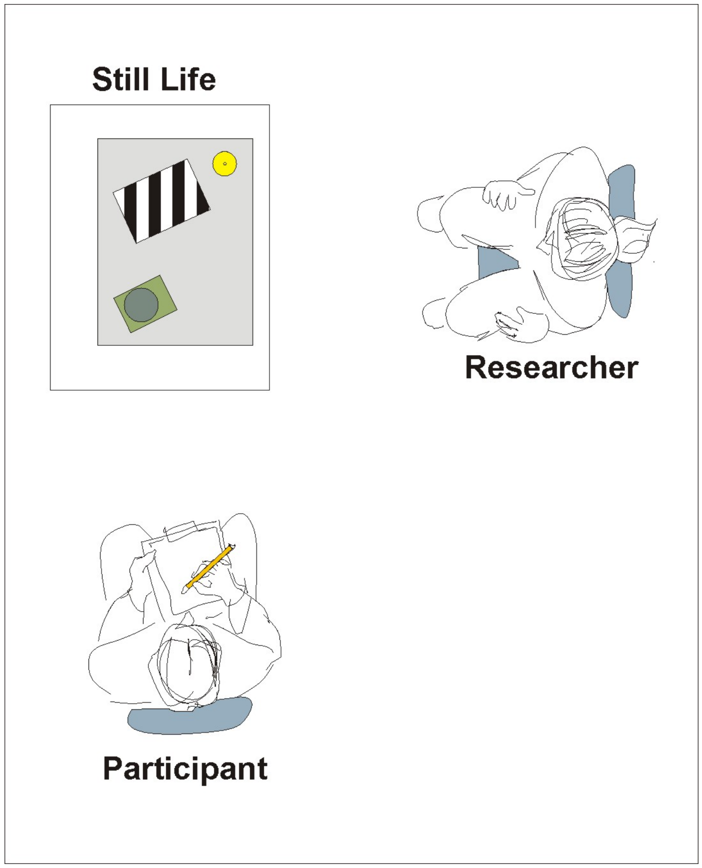

Figure 2.1: Overhead view of the experimental set-up.

Other aids We used a common protractor to review angle measures with participants, to measure the angles on the drawings and to measure the projected angles of selected edges on the ground truth photographs of the still lifes. We used a clinometer to measure the slant of selected edges in the still lifes. We constructed a simple planar angle display (Figure 2.3) with a pair of jointed arms that swiveled around a pop rivet, marked on the back with guidelines for reliably displaying the angles for a planar angle judgement task. We used a small toy ladder as a visual aid when explaining slant and projected angle judgement tasks.

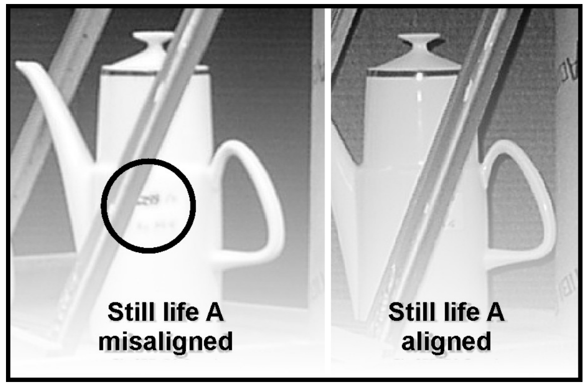

Figure 2.2: When the participant's viewpoint was correctly aligned with the still life, the small label on the teapot was occluded by the edge of the picture frame.

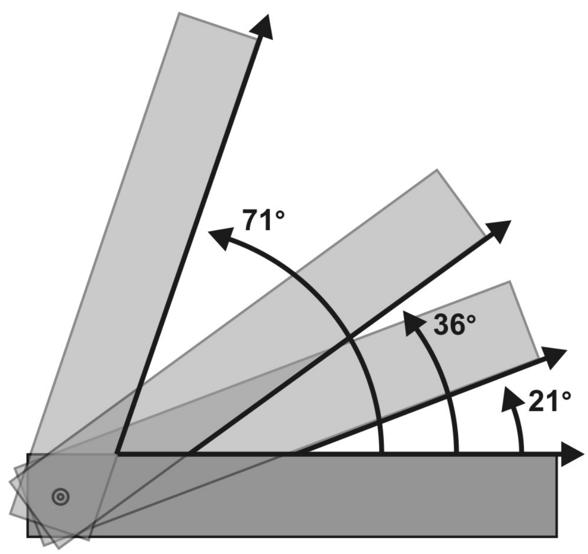

Figure 2.3: The display for the planar angle judgement task.

2.2.3 Procedure

Overview

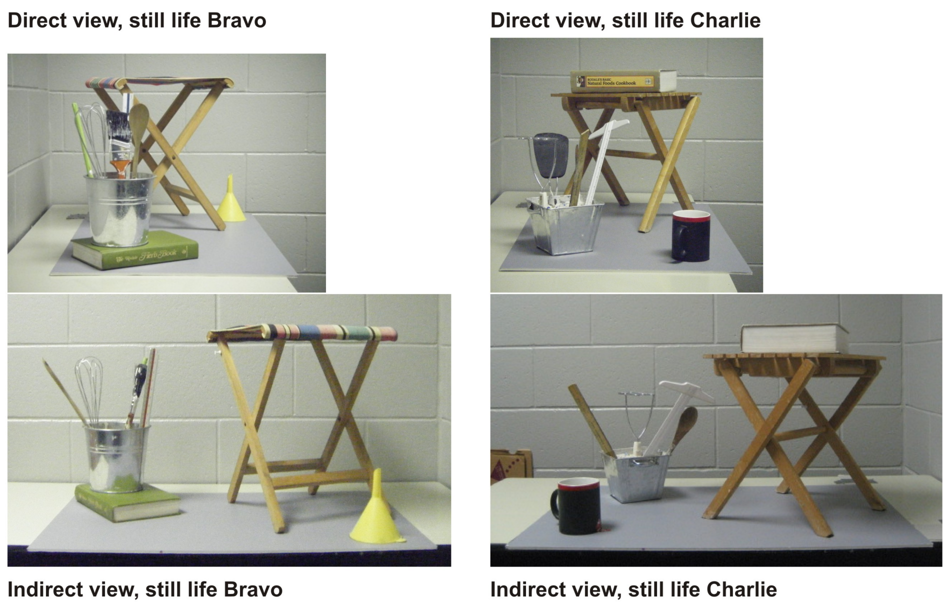

Figure 2.4: Left: Still life Bravo as it appeared to the participant, i.e., the "direct view" (above), and as it appeared to the researcher seated to the participant's right, i.e., the "indirect (side) view" (below).
Right: Still life Charlie as it appeared to the participant, i.e., the "direct view" (above), and as it appeared to the researcher seated to the participant's right, i.e., the "indirect (side) view" (below).

After giving informed consent, participants performed several tasks. Participants worked solo and it usually took about an hour for each of them to complete all of the study tasks. They completed a demographic survey, performed several perceptual tasks and made two drawings. In order to understand the perceptual and drawing tasks clearly, it is helpful to know that they were all performed from two different viewpoints, one real ("direct") and one imagined ("indirect" or "side"), on two different (but similar) still lifes. When a participant arrived for testing, still life Bravo was already assembled on a table (Figure 2.4, left). The participant sat on a chair, facing the still life from approximately two metres away. The researcher sat to the participant's right, facing the still life from a different viewpoint 90° away (Figure 2.1). She asked the participant to make several perceptual judgements (slant, projected angle, extent and occlusion; described below) about still life Bravo *as the participant imagined it would look when seen from the researcher's viewpoint.* We report this condition as the "indirect (side) view." Then the researcher asked the participant to draw still life Bravo from direct observation (i.e., as it looked from the participant's own viewpoint). We report this condition as the "direct view" and it was, at this point in the study, unremarkable to the participants. This is the way one would normally draw from a still life.

Then the researcher disassembled still life Bravo, replaced it with still life Charlie, and returned to her seat 90° away (Figure 2.4). She asked the participants to make the same perceptual judgements (slant, projected angle, extent and occlusion) about still life Charlie, but as it looked from their own viewpoint (i.e., the direct view). Finally, the researcher asked participants to draw still life Charlie, *not as they saw it but as they imagined it would look to her (i.e., the indirect (side) view).*

Fixed order of tasks For several reasons, the order of the tasks was fixed rather than counter-balanced. The drawing histories were assessed post hoc, so the researcher was blind to expertise level during the tasks. The tasks were carefully ordered to make the most of the participant's time by minimizing the need for still life set-up and change-over. In a previous still life drawing study (Carson & Allard, in press) where participants drew from one still life and made perceptual judgements from another, task order (drawing/judgement) had no effect on accuracy. Also in that study, even when given an unlimited amount of time in which to complete a drawing, some participants gave up on their drawing in five or ten minutes. Our primary concern in the present study was our expectation that drawing from an imagined indirect viewpoint would be discouragingly difficult and participants might give up on it. [1] We assigned the indirect drawing task last so that, if participants gave up because of it, we would already have data from the other, more important experimental tasks. We hoped to prevent the indirect drawing task from affecting performance in the other tasks, and to minimize attrition.

Every effort was made to design two similar still lifes in order to reduce variability: each still life included a book, a tin of tools, a cross-legged stool and a small object (one a mug, the other a funnel); and, each still life was positioned on a gray mat of the same size, occupying roughly the same footprint and volume. Participants made perceptual judgements from an indirect (side) view of still life Bravo, which should not affect their

[1] Researchers had previously reported that participants asked to draw a much simpler wire shape from an imagined indirect viewpoint said they did not feel as if they knew what that would look like (Rock et al., 1989).

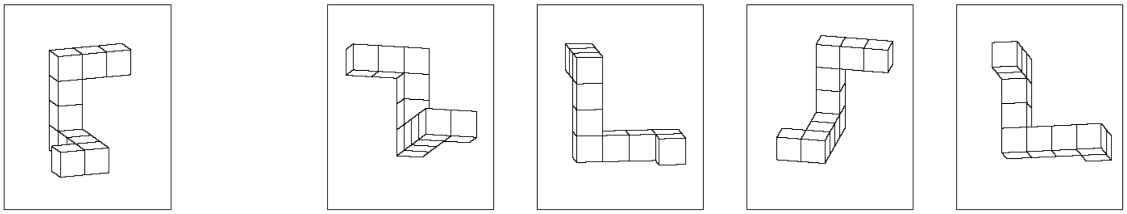

Figure 2.5: Sample question reproduced from the instructions of the Revised Vandenberg & Kuse Mental Rotations Test (Peters et al., 1995). It shows a target figure (left) and four stimulus figures (right), two of which (second and third) are rotated versions of the target.

drawing from a direct view. Similarly, they made perceptual judgements from a direct view of still life Charlie, which should not affect their drawing from an indirect (side) view.

Tasks

Survey Participants completed a short survey that captured basic demographic information as well as a history of their training and experience in art, drawing and drawing-related activities. We used the history of drawing experience as a self-report measure of expertise.

Planar angle judgement task In order to refresh participants' memory on the use of degrees as units of measure (and to detect any conspicuous misapprehensions), the researcher reviewed angles in degrees and demonstrated angles of 0° (horizontal), 30° and 90° (vertical). Then participants estimated the sizes of three different acute angles as shown on the planar angle display (Figure 2.3). The angles were arbitrarily chosen such that none was close to 0° or 90° or one another. The angles were marked clearly on the back of the planar angle display and presented in the same order for each participant: 36°, 71°, and 21°. Participants reported verbally their best estimates, in degrees, of each angle.

Mental rotation task Participants completed a 24-question, pencil-and-paper mental rotation task (Peters et al., 1995) identifying stimulus figures that are rotated versions of the target figure (Figure 2.5).

Perceptual judgement tasks from the indirect (side) view For the perceptual judgement tasks (and drawing tasks, described below), participants needed to maintain a consistent viewpoint on the still life. The researcher coached participants on how to align themselves such that the appropriate landmark was occluded. The participants were not fixed in position in any way, but encouraged to self-monitor their maintenance of the correct viewpoint, and to use the landmark to restore their position whenever they detected that they had strayed.

Slant judgement (indirect, Bravo) The researcher explained *slant*, which is the angle at which an edge or plane rises off a horizontal surface, such as the pitch of a roof or the angle at which a ladder leans against the wall. Participants were asked to imagine how still life Bravo looked from the researcher's viewpoint, and to report verbally their best estimates, in degrees, of the slants of three edges. For example, the researcher asked "What is the slant of the green bamboo stick, as seen from my position?" (Figure 2.6). The accurate slant measures (verified with a clinometer on the edges of still life Bravo) were

Figure 2.6: Participants were asked to estimate the slant of three edges in still life Bravo from an indirect (side) view, including the green bamboo stick (which had a slant of 71°). They were also asked to estimate the projected angle of three edges in still life Bravo from an indirect (side) view, including the wooden spoon (which had a projected angle of 55°).

71°, 81° and 76°. Because slant is a viewpoint-independent measure, estimating slant from the researcher's viewpoint should be exactly the same as estimating slant from one's own. The researcher did not mention this, nor did any participant volunteer that observation.

Projected angle judgement (indirect, Bravo) The researcher explained a *projected angle*, which is the angle the edge of an object would appear to make to the horizontal when projected onto the picture plane. That is, the projected angle is the angle at which one would draw that edge. When necessary, the researcher suggested the participant imagine looking through a window and tracing the edge of a distant subject onto the nearby glass. The angle that traced line makes to the horizontal is the projected angle of the edge. Then participants were asked to imagine how still life Bravo looked from the researcher's viewpoint, and to report verbally their best estimates, in degrees, of the projected angles of three edges. For example, the researcher asked "What is the projected angle of the wooden spoon, as seen from my position?" (Figure 2.6). The accurate projected angle measures (verified with a protractor on a ground truth photograph of still life Bravo taken from the researcher's viewpoint, Figure 2.4) were 55°, 73° and 73°.

Extent judgement (indirect, Bravo) Still imagining themselves sitting in the researcher's position, 90° away on their right, participants reported verbally on the extent of still life Bravo: its leftmost and rightmost points, its nearest and deepest points, and its highest and lowest points (on the picture plane). For example, "What is the nearest object in the still life, from my viewpoint?" to which the correct answer was the yellow funnel (Figure 2.4, lower left).

Occlusion judgement (indirect, Bravo) The participants answered four questions about whether particular objects in the still life occluded other objects, when seen from the researcher's position. For example, "Does the white ruler occlude any part of the stool, from my viewpoint?" to which the correct answer was "No." The four occlusion questions were chosen such that the correct answers were "No" to two and "Yes" to the other two (Figure 2.4, lower left).

Drawing task from the direct view Participants used pencil and paper (on a clipboard) to make a line drawing of still life Bravo (including the gray mat on which the still life stood), to the best of their ability, in ten minutes, from direct observation. The researcher encouraged them to draw in their usual way (which might include sketching, shading and erasing) but without the use of rulers or other straightedges, and then to emphasize final lines, if necessary, to make them unambiguous.

Perceptual judgement tasks from the direct view After the drawing task was completed, the researcher removed still life Bravo. She assembled still life Charlie in its place and coached the participants on how to align themselves with it consistently. They performed the same four perceptual judgement tasks, but from a direct view of still life Charlie: slant judgement (accurate measures were 68°, 56° and 52°), projected angle judgement (accurate measures were 77°, 75° and 70°), extent judgements, and occlusion judgements. Many participants mentioned that the extent and occlusion questions were easy from direct observation.

Drawing task from the indirect (side) view Participants used pencil and paper (on a clipboard) to make a line drawing of still life Charlie (including the gray mat on which the still life stood), to the best of their ability, in ten minutes, *as they imagined it looked from the researcher's indirect (side) viewpoint seated to the right of the still life,* 90° *away.*

2.2.4 Measuring angle error

We collected participants' angle judgements as verbal reports. We measured angle drawing accuracy directly on each drawing with a protractor. We took the size, in degrees, of every independent non-trivial angle present in the drawing for comparison to the same angles on a ground truth photograph. By "non-trivial" angle we mean angles salient enough to be present and practical to measure in some or most drawings. We omitted angles of very short edges and angles close to 180°. By "independent" angle we mean any angle whose size was geometrically independent of all the other angles measured. We omitted those angles which complemented other angles already measured in order to be consistent and not to bias results. For example, when two edges cross to form an "X" shape, there are four possible angles to measure, each of which would have an error of identical magnitude. For another example, when three edges meet to form a triangle, there are three interior angles to measure, but since they sum to 180°, only two of those angles can be independent (and *any two* of those angles can be independent).

All the angle errors—drawing and judgement—are reported here as the magnitude of the difference between the reported angle and the correct value, scaled to the percentage of the correct value.

2.3 Results

2.3.1 Drawing experience and angle accuracy

We measured 28 independent non-trivial angles in the drawings made from a direct view of still life Bravo. For each participant, we defined accuracy—a performance-based measure of

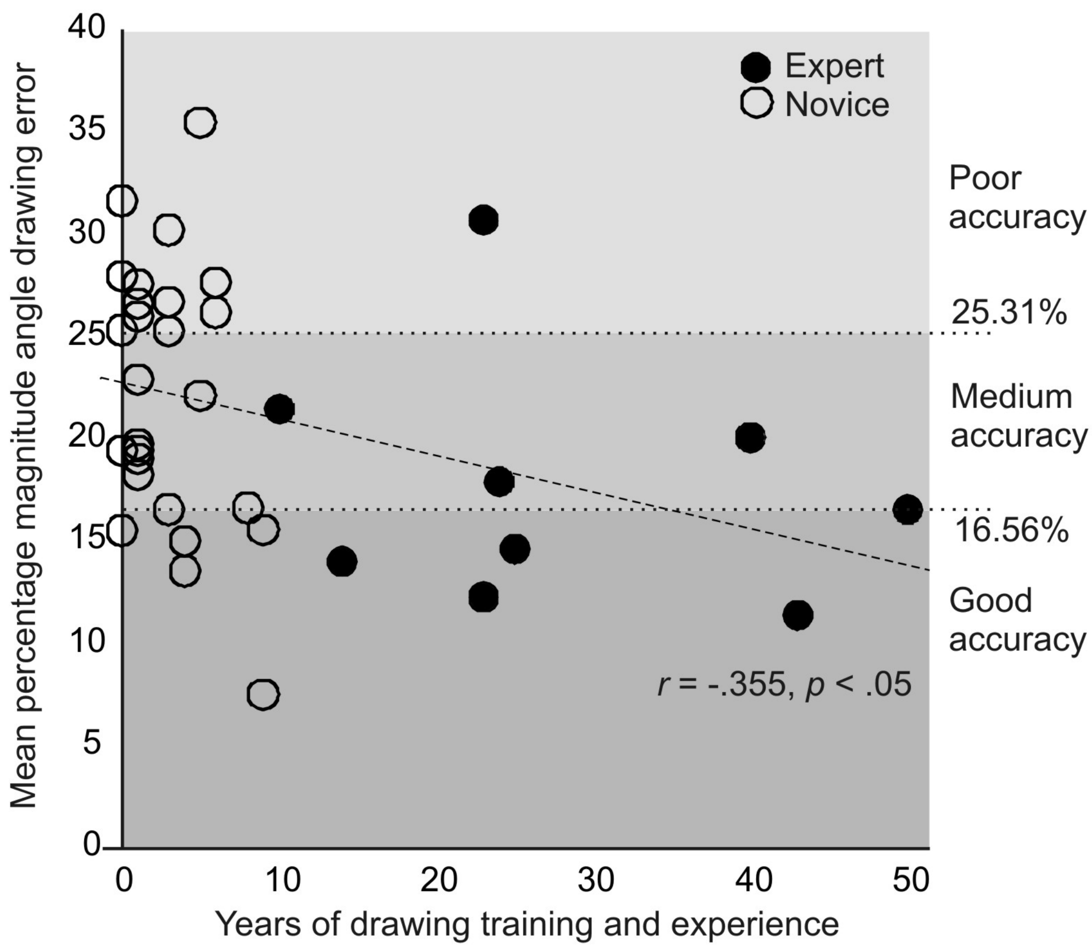

Figure 2.7: Drawing accuracy as a function of years of art experience.

drawing skill—as the mean of the percentage magnitude angle drawing error across all the angles present in the drawing. Experts with ten more more years of experience were more accurate than novices, a difference that approached significance ($t(33) = .686$, $p = .071$). Accuracy correlated well with participants' years of art experience ($r(35) = -.355$, $p < .05$; Figure 2.7).

For analysis, we stratified the participants into groups of (nearly) uniform size on three levels of drawing accuracy, good ($N = 11$), medium ($N = 12$) and poor ($N = 12$). We also divided participants into two groups on the basis of their years of drawing experience, experts ($N = 9$) and novices ($N = 26$), using the conventional ten-year threshold. We analyzed performance on the experimental tasks, where appropriate, with respect to both drawing *experience* (expert/novice; a self-report measure) and drawing *accuracy* (good/medium/poor; a performance-based measure).

2.3.2 Planar angle judgement, angle accuracy and drawing experience

There was no significant difference in planar angle judgement between the three groups of drawing accuracy, $F(2,31) = 0.505$, $p = .609$, $\eta^2 = .032$, nor between experts and novices, $F(1,31) = 2.603$, $p = .117$, $\eta^2 = .077$.

2.3.3 Local drawing accuracy within the image

Local drawing accuracy

By inspection it was clear that accuracy was not uniform across the image (Figure 2.8). Some angles were clearly more difficult to draw than others. On an angle-by-angle basis, experts were usually—but not always—more accurate than novices. Participants with good accuracy were generally more accurate, angle by angle, than those with medium accuracy, who were generally more accurate angle by angle than those with poor accuracy (Figure 2.9).

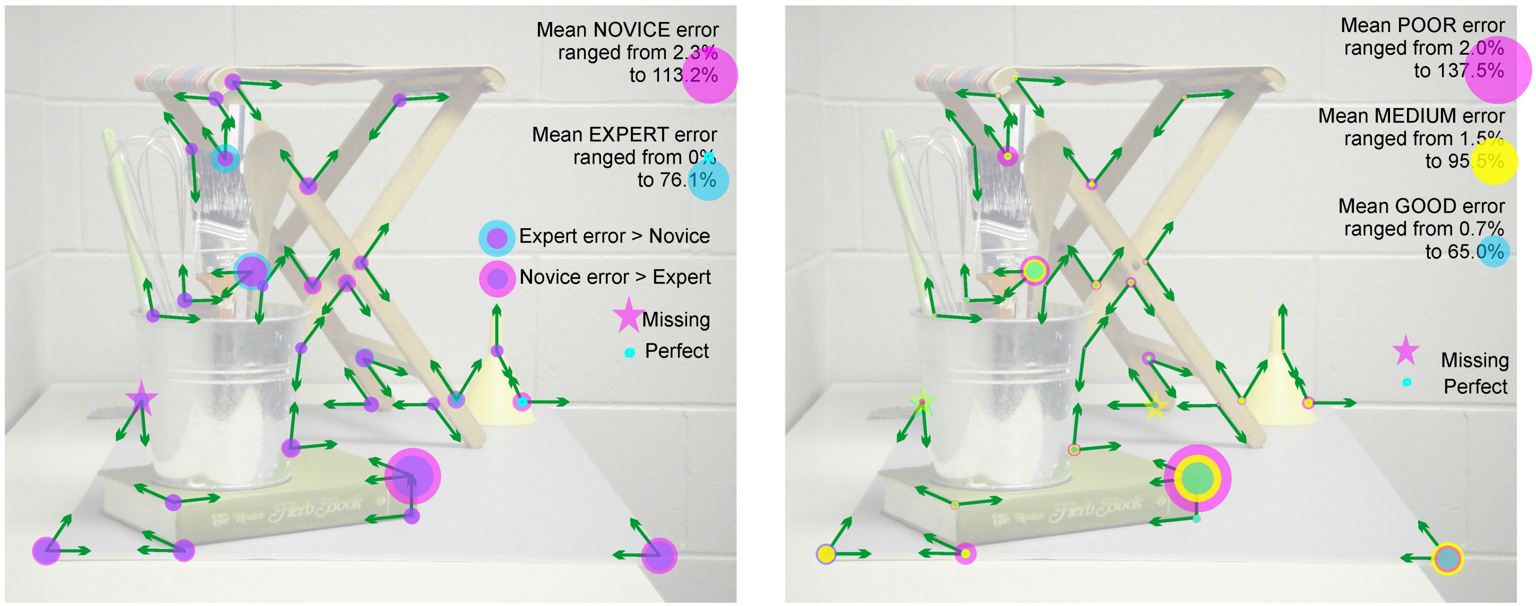

Figure 2.8: The arrows show the location of every angle measured on the ground truth photograph of still life Bravo. Superimposed on each vertex are circles whose diameter are proportional to the mean error on that angle.

Left: The magenta circle represents the mean error of novice participants (less than ten years of drawing experience). The cyan circle represents the mean error of expert participants (ten or more years of drawing experience). Most circles form a bullseye with a purple centre and magenta surround indicating that mean novice error was greater than expert on that angle. A few circles have a cyan surround, highlighting angles where mean expert error was greater than novices.

Right: The cyan circle represents the mean error for *good* accuracy. The yellow circle represents the mean error for *medium* accuracy. The magenta circle represents the mean error for *poor* accuracy. The circles form a bullseye with a dark centre. The colour of the outermost ring reveals which group had the largest error on that particular angle. By observation, that is not always the participants with *poor* accuracy overall.

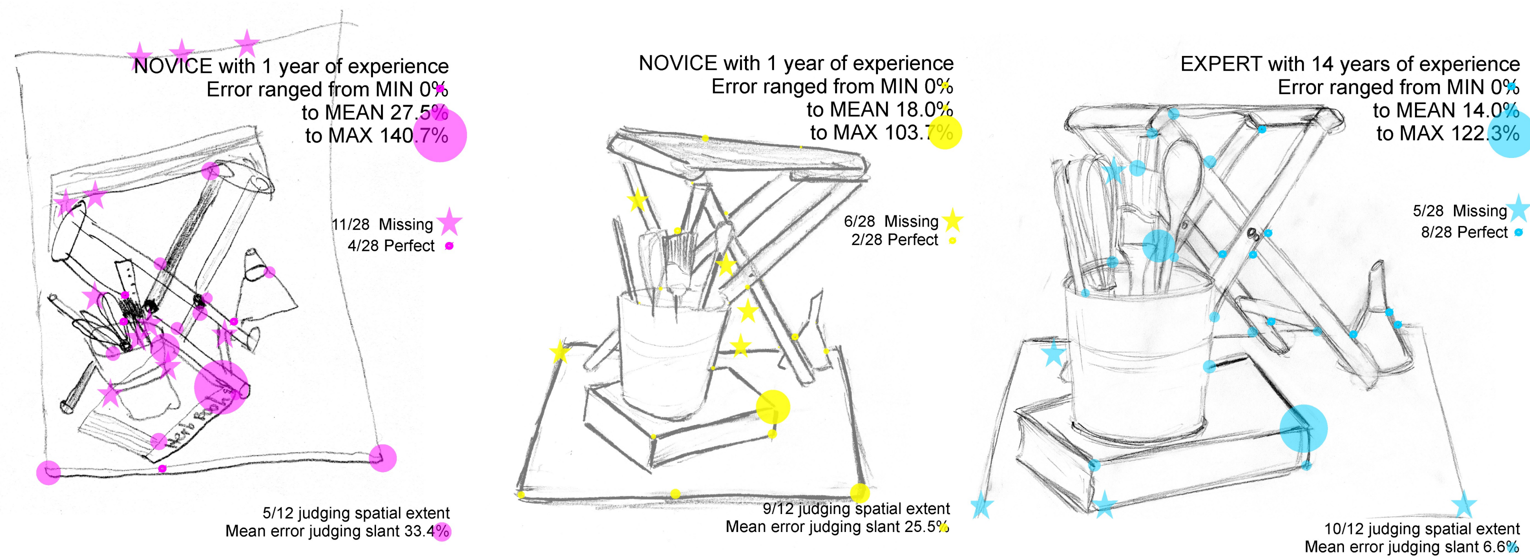

Figure 2.9: Left: This drawing was made by a participant with 1 year of drawing experience, whose performance placed her in the *poor* level of accuracy. The magenta circles represent the size of the error at that point, which ranged from zero to 140.7%, with a mean of 27.5%. The participant missed eleven angles (marked here with magenta stars) and drew four angles perfectly (marked here with ϕ).
Middle: This drawing was made by a participant with one year of drawing experience, whose performance placed him in the *medium* level of accuracy. The yellow circles represent the size of the error at that point, which ranged from zero to 103.7%, with a mean of 18.0%. The participant missed six angles (marked here with yellow stars) and drew two angles perfectly (marked here with ϕ).
Right: This drawing was made by a participant with 14 years of drawing experience, whose performance placed her in the *good* level of accuracy. The cyan circles represent the size of the error at that point, which ranged from zero to 122.3%, with a mean of 14.0%. The participant missed five angles (marked here with cyan stars) and drew eight angles perfectly (marked here with ϕ).

Acute and obtuse angles

Sixteen of the angles measured in the drawings of still life Bravo were acute and 11 were obtuse. We compared the mean percentage magnitude angle drawing error for acute and obtuse angles, which differed significantly ($F(1, 31) = 404.221$, $p < .001$, $\eta^2 = .929$; Figure 2.10, left). People were more accurate on obtuse angles than acute.

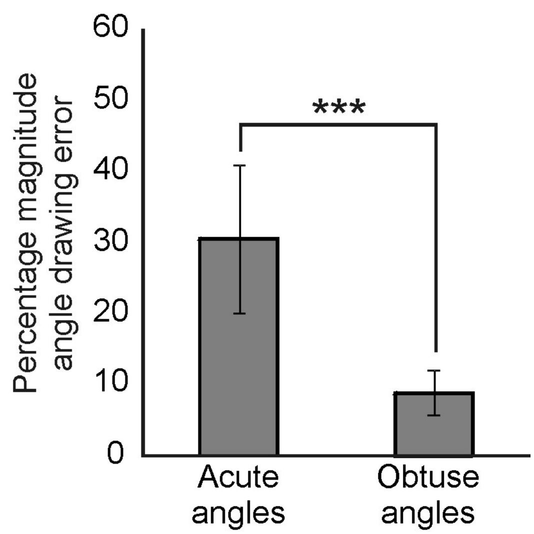

Figure 2.10: Drawing error as a function of angle size, acute and obtuse. Error bars represent ± one standard deviation.

There was no significant interaction with level of drawing experience, ($F(1, 31) = 1.117$, $p = .299$, $\eta^2 = .035$).

Divergent perspective

The angle with the greatest mean percentage magnitude angle drawing error was the rightmost corner of the book cover (ground truth 27°; error M= 102.72%, SD= 52.51; Figure 2.8). It was drawn much larger than its true size throughout the dataset, giving depictions of the book a characteristic divergent perspective. This was also the smallest angle we measured. To test whether this error was entirely due to angle size, we compared it with another small angle nearby (ground truth 30°; error M= 27.92%, SD= 27.14; Figure 2.11). Drawing error on these two angles differed significantly ($F(1, 19) = 35.157$, $p< .001$, η^2= .649; Figure 2.12, left).

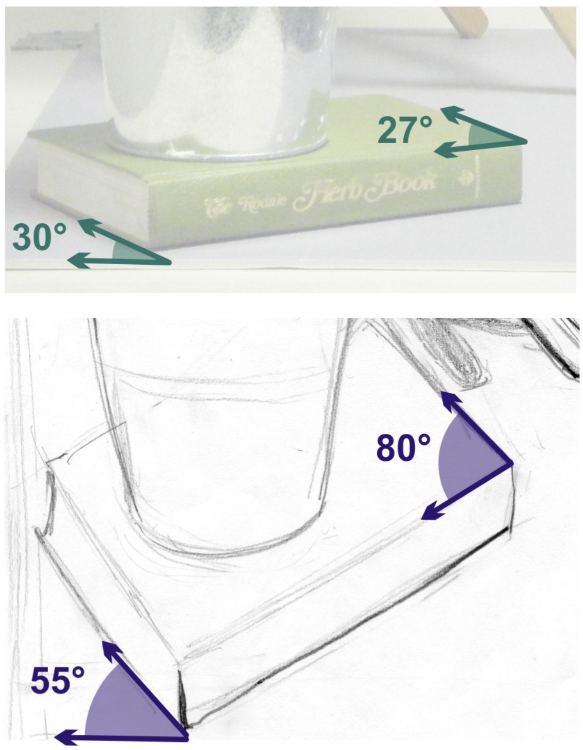

Figure 2.11: Detail views of the two angles of interest marked on the ground truth photograph of still life Bravo (above) and a participant's drawing (below).

There was no significant interaction between accuracy on these two small angles and experience ($F(1, 19) = .081$, p= .779, η^2= .004).

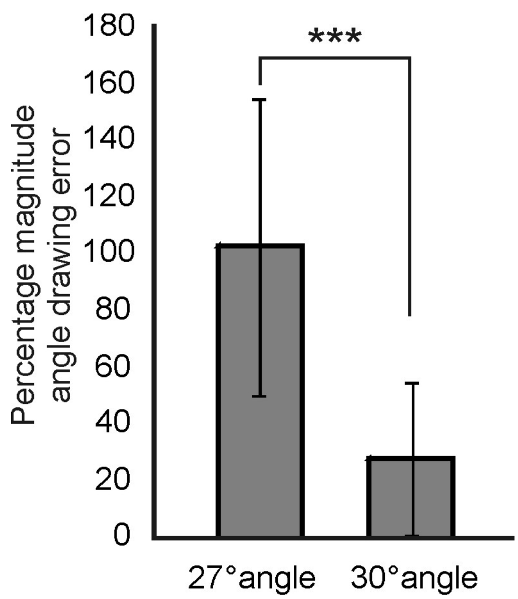

Figure 2.12: Drawing error as a function of angle size, 27° (the book in Figure 2.11) and 30°. Error bars represent ± one standard deviation.

2.3.4 Drawing experts and a 3D model of space

We used a mental rotation test and several novel tasks to probe the participants' understanding of the three-dimensional properties of the still life. We summarize the findings in Table 2.1, and report them in detail below.

Evidence of a better 3D model with drawing experience & accuracy		
Task	Experience	Accuracy
Mental rotation	*n.s.*	*n.s.*
Projected angle judgement	$p = .091$	*n.s.*
Slant judgement	*n.s.*	$p < .05$
Occlusion judgement	*n.s.*	*n.s.*
Extent judgement	$p < .05$	*n.s.*
Missing angles when drawing, indirect	*n.s.*	$p = .095$
Accuracy when drawing, indirect	*n.s.*	*n.s.*

Table 2.1: Summarizing the evidence, by task, that experience and accuracy have significant relationships with a participant's understanding of the three-dimensional properties of the still life.

Mental rotation

There was no significant relationship between participants' mental rotation score and their drawing experience, $F(1, 29) = .034$, $p = .855$, $\eta^2 = .001$, nor between mental rotation and level of drawing accuracy, $F(2, 29) = .020$, $p = .980$, $\eta^2 = .001$.

Projected angle judgement, direct and indirect

We used a repeated-measures ANOVA to compare projected angle judgement from direct and indirect (side) views, by participant accuracy and experience. Participants were significantly more accurate in making projected angle judgements from the direct view (error $M = 13.94\%$, $SD = 10.07$) than from the indirect (side) view (error $M = 21.82\%$, $SD = 10.97$); $F(1, 31) = 11.356$, $p < .01$, $\eta^2 = .268$ (Figure 2.13, left). There was no significant main effect of accuracy level on projected angle judgement, $F(2, 31) = 1.390$, $p = .264$, $\eta^2 = .082$ (Figure 2.13, centre). However, there was a trend toward significance for the effect of drawing experience on projected angle judgement, $F(1, 31) = 3.051$, $p = .091$, $\eta^2 = .090$

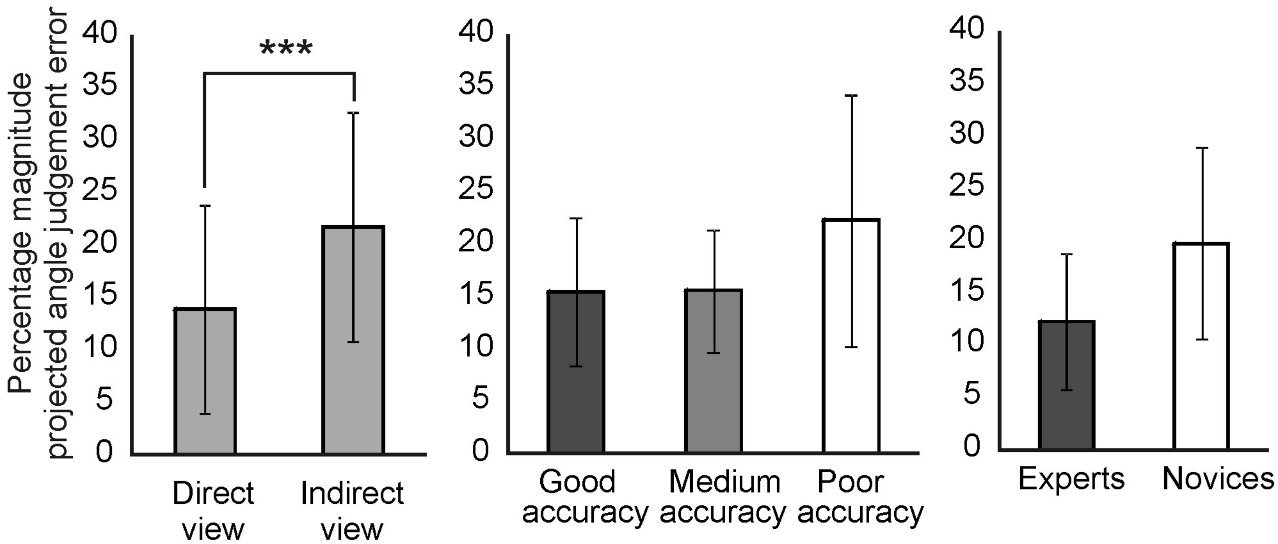

Figure 2.13: Projected angle judgement error as a function of direct or indirect view (left), as a function of participants' level of accuracy (centre), and as a function of participants' level of experience (right). Error bars represent ± one standard deviation.

(Figure 2.13, right). Experts made smaller errors (M= 12.27%, SD= 6.51) than novices (M= 19.82%, SD= 9.21) when judging projected angles from the still life.

Slant judgement, direct and indirect

We used a repeated-measures ANOVA to compare slant judgement from direct and indirect (side) views, by participant accuracy and experience. Unexpectedly, participants were significantly more accurate making slant judgements from the *indirect* (side) view (error M= 9.91%, SD= 4.36) than from the direct view (error M= 20.25%, SD= 9.39), $F(1, 31) = 39.100$, $p<$ .001, η^2= .558 (Figure 2.14, left). Participants' overall level of drawing accuracy was significantly related to their slant judgement, $F(2, 31) = 4.210$, $p<$ .05, η^2= .214 (Figure 2.14, centre). Participants with good drawing accuracy made smaller errors of slant judgement (M= 11.36%, SD= 3.20) than those with medium accuracy (M= 17.41%, SD= 4.52) or poor accuracy (M= 16.16%, SD= 6.27). A post hoc comparison between levels of angle accuracy, with Bonferroni correction, revealed that the difference between the good and medium accuracy groups was significant ($p<$ .05), the difference between the good and poor accuracy groups trended toward significance (p= .081), and there was no meaningful difference between the slant judgements of the medium and

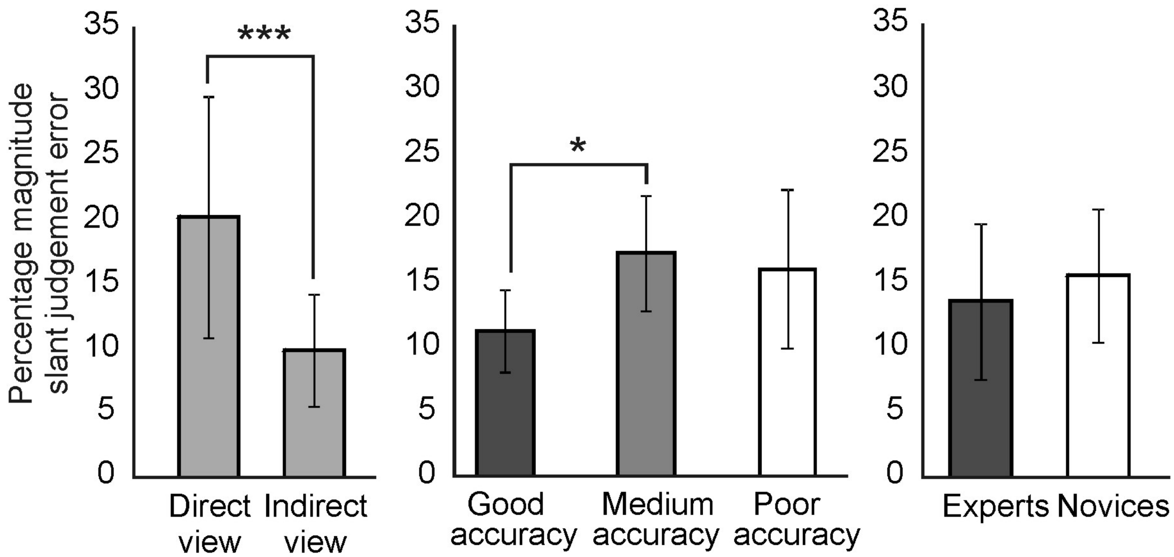

Figure 2.14: Slant judgement error as a function of direct or indirect view (left), as a function of participants' level of accuracy (centre), and as a function of participants' level of experience (right). Error bars represent ± one standard deviation.

poor groups (p= 1.00). The interaction between accuracy and viewpoint tended toward significance ($F(2, 31) = 2.713$, p= .082, η^2= .149).

There was no significant relationship between drawing experience and slant judgement $F(1, 31) = .052$, p= .822, η^2= .002 (Figure 2.14, right), nor any interaction with viewpoint.

Occlusion judgement, direct and indirect

All participants scored perfectly when making occlusion judgements from the direct view (number of correct answers $M = 4.00$, $SD = 0.00$). Performance was less accurate on the indirect (side) view (number of correct answers $M = 3.63$, $SD = 0.598$). Comparing occlusion judgement in the direct and indirect views with a Wilcoxon signed-rank test, the difference was significant ($z = -3.127$, $p < .01$, $r = -.370$; (Figure 2.15). There was

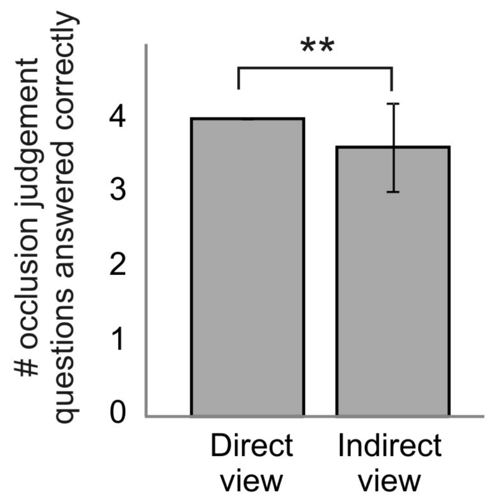

Figure 2.15: Occlusion judgement accuracy as a function of direct and indirect view. All participants answered all four occlusion judgement questions correctly in the direct view condition. Error bars represent ± one standard deviation.

no significant difference in occlusion judgement in the indirect condition by experience ($F(1, 31) = .033$, $p = .856$, $\eta^2 = .001$) or accuracy ($F(2, 31) = 1.380$, $p = .267$, $\eta^2 = .082$).

Extent judgement, direct and indirect

When we used a repeated-measures ANOVA to compare extent judgement from direct and indirect (side) views, by participant accuracy and experience, we found that there was no main effect of viewpoint ($F(1, 31) = .327$, $p = .572$, $\eta^2 = .010$). This was unique amongst the direct/indirect perceptual judgement tasks. All other tasks (projected angle, slant and occlusion judgement) showed a significant main effect of viewpoint.

Experience is significantly related to extent judgement, $F(1, 31) = 5.478$, $p < .05$, $\eta^2 = .150$ (Figure 2.16). Experts made more correct extent judgements from the direct viewpoint

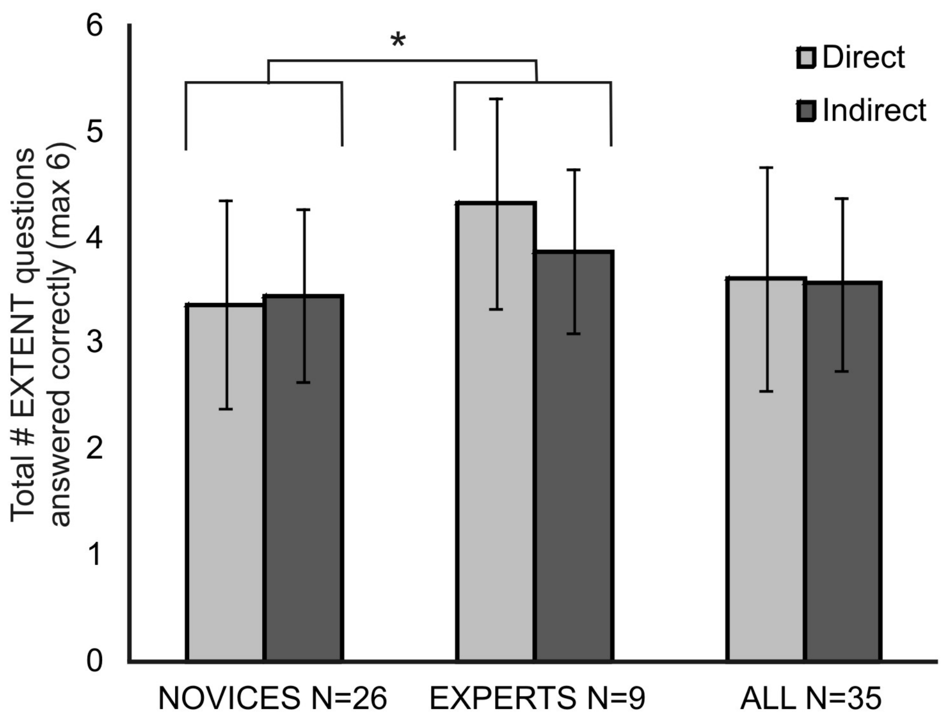

Figure 2.16: Extent judgement accuracy as a function of direct and indirect view, grouped by participants' level of drawing experience. Error bars represent ± one standard deviation.

(M= 4.33/6, SD= 1.00) than did novices (M= 3.38/6, SD= 0.98), and made more correct extent judgements from the indirect viewpoint (M= 3.88/6, SD= 0.78) than did novices (M= 3.46/6, SD= 0.81).

There was no significant relationship between levels of accuracy and extent judgement ($F(2, 31) = .814$, p= .452, η^2= .050).

2.3.5 Drawing from an imagined, indirect (side) viewpoint

We measured 25 independent non-trivial angles on the drawings made from an imagined, indirect (side) view of still life Charlie. For each participant, we defined *indirect accuracy* as the mean of the percentage magnitude angle drawing error across all the angles present in the drawing.

Missing angles

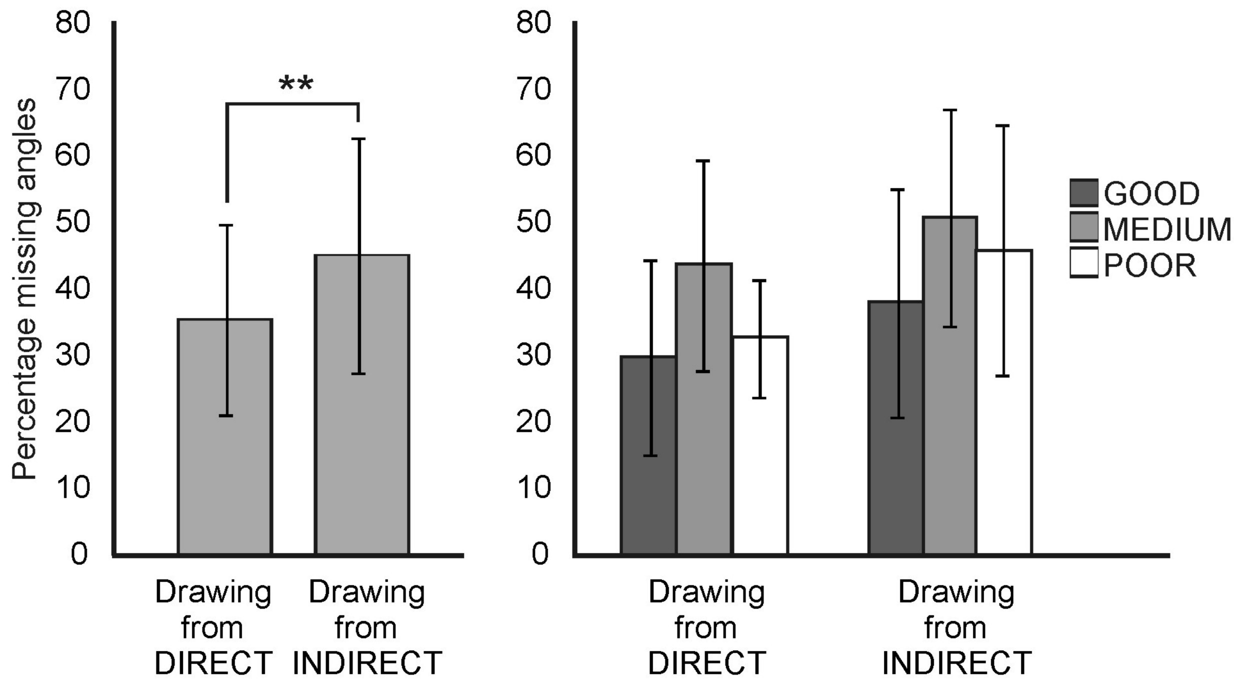

Figure 2.17: Percentage of missing angles as a function of direct or indirect view (left) and grouped by participants' level of drawing accuracy (right). Error bars represent ± one standard deviation.

Incomplete drawings had many missing angles. Missing angles were also symptomatic of drawings with serious structural flaws. In such drawings, edges which—in the ground truth—should cross or occlude others to form angles did not overlap at all. Drawings of still life Charlie missed more angles than drawings of still life Bravo. A repeated-measures ANOVA comparing the percentage of missing angles between the direct and indirect viewpoints showed a significant main effect of viewpoint, $F(1,31) = 10.162$, $p < .01$, $\eta^2 = .247$ (Figure 2.17, left). When participants drew from a direct viewpoint of still life Bravo, they omitted fewer angles (missing $M = 35.61\%$, $SD = 14.45$) than when they drew from an indirect (side) viewpoint of still life Charlie (missing $M = 45.25\%$, $SD = 17.80$).

There was a trend toward a relationship between overall accuracy levels (when drawing from the direct viewpoint) and missing angles $F(2,31) = 2.540$, $p = .095$, $\eta^2 = .141$ (Figure 2.17, right). Participants with good accuracy (missing $M = 33.79\%$, $SD = 14.00$) omitted a lower percentage of angles from those with medium accuracy (missing $M = 47.16\%$, $SD = 15.00$), a difference that trended toward significance ($p = .089$) but neither of those groups was distinguishable from those with poor accuracy (missing $M = 38.99\%$, $SD = 11.88$).

There was no significant relationship between experience and missing angles $F(1,31) = .034$, $p = .855$, $\eta^2 = .001$.

Indirect accuracy

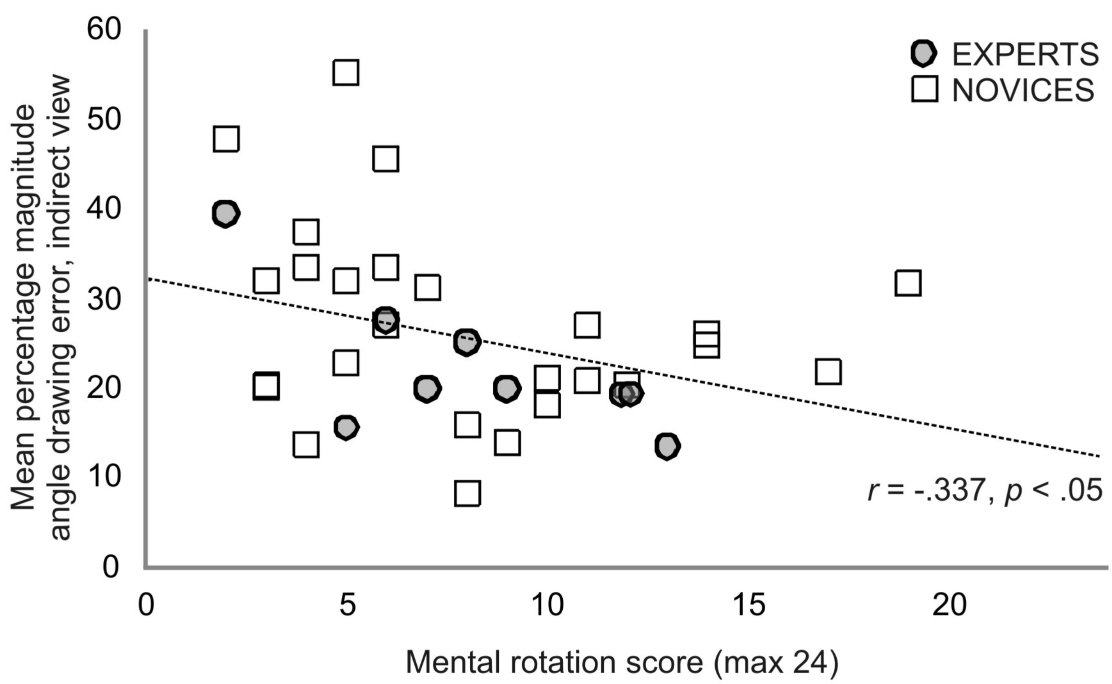

Figure 2.18: Angle drawing error *when drawing from the indirect view* as a function of mental rotation task score.

Indirect accuracy correlates with the mental rotation task, $r(35) = -.337$, $p < .05$ (Figure 2.18).

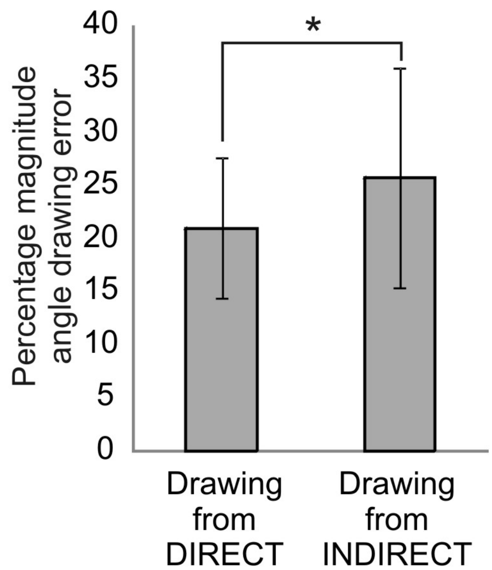

Figure 2.19: Drawing error as a function of direct and indirect view. Error bars represent ± one standard deviation.

A paired-samples t-test showed that participants' overall accuracy differed significantly between the direct and indirect (side) viewpoints, $t(34) = 2.450$, $p < .05$ (Figure 2.19). Participants had higher error from the indirect (side) view ($M = 25.81\%$, $SD = 10.39$) than from the direct view ($M = 21.08\%$, $SD = 6.59$).

A one-way ANOVA on indirect accuracy showed no significant relationship with drawing accuracy in the direct condition ($F(2, 31) = .309$, $p = .737$, $\eta^2 = .020$) or experience ($F(1, 31) = .822$, $p = .372$, $\eta^2 = .026$).

Characterizing indirect drawings by observation

As the basis for future work, we collected some observations based not on angles but on possible patterns in the depiction and arrangement of the main objects in the still life. Our sample size was too small for quantitative analysis so we report the qualitative results here in simple tables.

Bird's eye view We observed that drawing the cross-legged stool, challenging enough from a direct view, was very difficult from an indirect (side) view for most of the participants. This was not a case of people making mistakes of which they were unaware; many participants complained about the indirect drawing task while they were doing it, and some singled out the stool as a particular problem. We can see, on one participant's drawing (Figure 2.20), a series of small sketches along the top, evidence of efforts to work out the relationship of the legs and seat from the researcher's viewpoint.

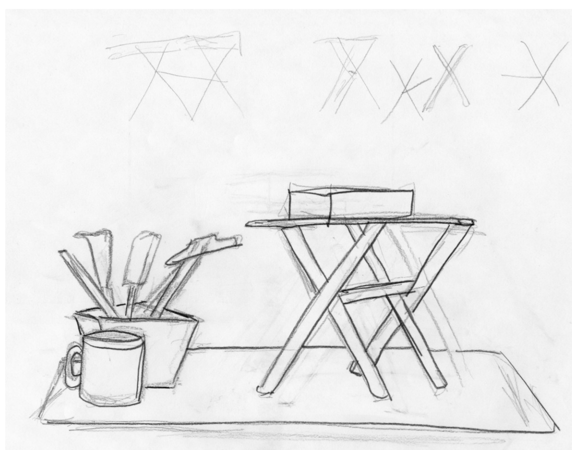

Figure 2.20: This participant, who had nine years of drawing experience and a mental rotation score of 8/24, made a series of small working drawings along the top of her page to figure out what the arrangement of the legs and seat of the stool should be when seen from the researcher's indirect (side) view. We can also see from the erasures that she drew and then revised her depiction of the stool in the finished drawing. She correctly depicted the seat of the stool at eye level. The mug was fairly well placed left/right and near/deep. (Compare with Figure 2.4.)

To describe the relative success of the drawing of the stool, we noted the depiction of its seat. Most people made a characteristic perspective error, showing the seat of the stool as it would appear from somewhat (or entirely) overhead rather than at eye level. We inspected the drawings of the seat and categorized each as eye level (which would correctly match the ground truth photograph), somewhat below eye level, a bird's eye view, or missing (Figure 2.21).

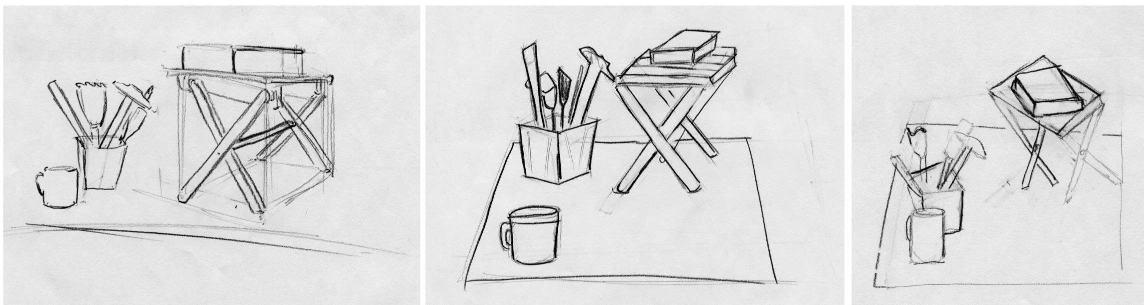

Figure 2.21: Left: This participant, who had 24 years of drawing experience and a mental rotation task score of 13/24, correctly depicted the seat of the stool as if it were at eye level, offering no view of its top face. This was the only drawing to correctly show the mug entirely to the left of the tin of tools. (Compare with Figure 2.4.)
Centre: This participant, who had one year of drawing experience and a mental rotation task score of 11/24, depicted the seat of the stool as if seen from somewhat overhead, which is incorrect. The mug was fairly well placed left/right but too near the viewer relative to the tin of tools.
Right: This participant, who had three years of drawing experience and a mental rotation task score of 4/24, depicted the seat of the stool from a bird's eye view, which is incorrect and is inconsistent with the view of the coffee cup, lower left. The mug was fairly well placed left/right.

Five novices—and no experts—made the bird's eye view error. People who drew the direct view with good accuracy never made the bird's eye view error in the indirect drawing (Table 2.2).

Depiction of stool seat when drawing from imagined indirect viewpoint					
	Experience			Accuracy	
	Novice	Expert	Poor	Medium	Good
	N= 26	N= 9	N= 12	N= 12	N= 11
Eye-level (correct)	30.7%	44.4%	33.3%	41.6%	27.2%
Below eye-level	38.4%	55.5%	41.6%	25.0%	63.6%
Bird's eye view	19.2%	0%	16.6%	25.0%	0%
Missing	11.5%	0%	8.3%	8.3%	9.0%

Table 2.2: Summary of how many participants, when drawing still life Charlie from an imagined indirect (side) view, drew the seat of the stool correctly as if at eye level, less correctly as if somewhat below eye level, or, incorrectly, from a bird's eye view.

Visualising the mug Four participants did not draw the handle of the mug, perhaps because they forgot or because they mistakenly believed the handle was hidden from the researcher's viewpoint. Everyone else drew the handle, correctly, on the left side of the mug. However, very few participants were successful at locating the mug in space relative to the stool and the tin of tools (Table 2.3). In the indirect (side) view of still life Charlie, the mug sits to the left of the tin of tools without occluding it, and only slightly nearer to the viewer. (Compare with Figure 2.4.) Only one participant correctly placed the mug to the left of the tin of tools (Figure 2.21, left). The rest erred, to varying degrees, in their left/right and near/far placement of the mug relative to the tin of tools (Figure 2.22). Two participants placed the mug *completely* to the right of the tin of tools, revealing a poor mental model of the 3D space (Figure 2.22, right).

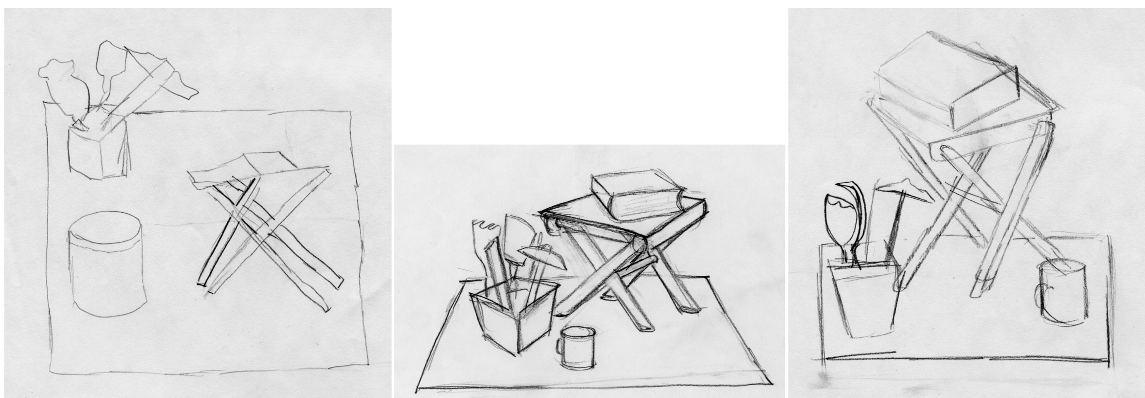

Figure 2.22: Left: This participant, who had three years of drawing experience and a mental rotation task score of 3/24, drew the mug incorrectly. Either it has no handle or the handle is incorrectly occluded by the body of the mug. It was fairly well placed left/right but much nearer the viewer (relative to the tin of tools) than it should be.
Centre: This participant, who had ten years of drawing experience and a mental rotation task score of 12/24, positioned the mug too far right but fairly accurately near/far relative to the tin of tools.
Right: This participant, who had three years of drawing experience and a mental rotation task score of 5/24, drew the mug far to the right of the tin of tools, even to the right of the stool. The mug and the tin of tools were shown, incorrectly, as being equally near/far from the viewer, but this would be better described as an error in the placement of the tin of tools, which was incorrectly shown nearer to the viewer than the stool. Note the use of divergent perspective on the mat.

Position of mug relative to tin of tools when drawing from imagined indirect viewpoint,					
	Experience		Accuracy		
	Novice	Expert	Poor	Medium	Good
	N= 26	N= 9	N= 12	N= 12	N= 11
Left (correct)	0%	11.1%	0%	8.3%	0%
Left & overlapping	50.0%	44.4%	58.3%	50.0%	36.3%
Overlapping	38.4%	22.2%	25.0%	25.0%	54.5%
Right & overlapping	7.9%	11.1%	8.3%	8.3%	9.0%
Right	3.8%	11.1%	8.3%	8.3%	0%

Table 2.3: Summary of how many participants, when drawing still life Charlie from an imagined indirect (side) view, drew the mug in its correct position to the left of the tin of tools, or in positions further to the right.

2.4 Discussion

For this study, we used a mental rotation test and a small battery of novel tasks to probe people's mental models of 3D space. Mean percentage magnitude angle drawing error is a satisfactory accuracy measure: practicable, statistically powerful, objective and informative about drawing properties both local and global. It is a continuous performance-based measure which satisfactorily captures expertise. People who draw more accurately or have more drawing experience have a better mental model of 3D space than novices.

Our results provide further evidence that context biases drawing outcomes. Tchalenko has shown, for example, that people are well able to copy lines in isolation that they struggle with when seen as part of a more complex image (2009). Mitchell and colleagues (2005) showed that people copying a drawing of the famous Shepard tables made larger errors than people drawing parallelograms of the same shape (but without the table legs that created the illusion; Figure 2.23). In our present study, we saw that a small angle in the bottom left of the image was drawn quite accurately, but a nearby angle of similar magnitude, part of the receding face of a rectangular prism (the book), was the worst drawn angle in the entire image. As researchers have observed before, there is a top-down influence of knowledge on perception and performance (Taylor & Mitchell, 1997). Drawing teachers have observed and lamented the same effect, and many exercises have been devised to illustrate or perhaps remedy it. For example, students are asked to draw the spaces between objects (the so-called "negative space" (Gombrich, 1960/1972)), to draw without looking at the page ("blind contour line drawing" (Nicolaides, 1941)), and

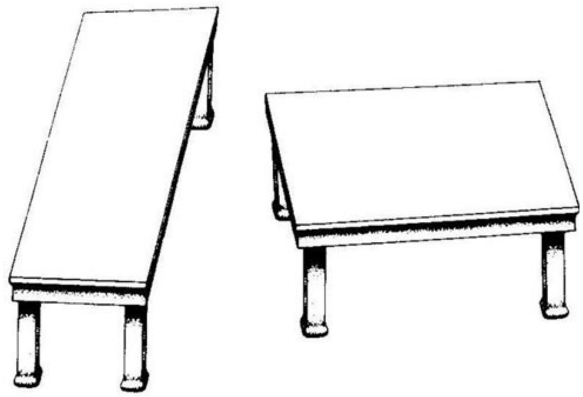

Figure 2.23: The famous tabletop illusion from Shepard (1990) demonstrates the powerful impact of context on perception. First, the two tabletops are formed by parallelograms, but the sides appear to diverge as the tables recede away from us in pictorial space. Second, and even more incredibly, the two tabletops are congruent parallelograms.

to draw from an inverted subject (Edwards, 1979). One of the basic motivations for this study and this thesis is laying the groundwork to investigate the usefulness of such exercises and to develop evidence-based drawing tuition. For example, we're encouraged by the recent finding of Cohen and Earls (2010) that the traditional assignment to draw from an upside-down image, no matter how persuasive it may be in the classroom, does not actually produce drawings that are measurably more accurate. The next chapter of this thesis compares drawing performance in the positive and negative space. We want to validate classic methods, when possible, and test new hypothesis-driven methods for teaching, learning and practising drawing better.

Mental rotation abilities were related to drawing in complex ways. While no relationship was found with direct accuracy, for the indirect condition (which more directly draws upon mental rotation abilities) there was a clear relationship between accuracy and mental rotation scores. However, the literature on mental rotation suggests that the demographics of our participant pool would confound or obscure any such relationship. Men are reported to have a strong advantage over women in mental rotation (Jansen & Heil, 2010; Linn & Petersen, 1985; Voyer, Voyer, & Bryden, 1995); mental rotation ability has been shown to diminish with age (Dror & Kosslyn, 1994); and, mental rotation ability is particularly strong in people in the science, technology, engineering and math (STEM) disciplines (Uttal & Cohen, 2012). Since we determined expertise post hoc and accuracy based on performance, we could not balance the expert/novice participants or participants of good/medium/poor accuracy on sex, age or participation in the STEM disciplines. Even in our unbalanced and heterogeneous sample (Appendix C), there was a significant main effect of sex ($F(1, 33) = 5.815$, $p < .05$, $\eta^2 = .150$), but we did not have the statistical power to simultaneously measure the effects of and interactions between sex, age, STEM education, expertise and accuracy.

One of the key findings supporting our hypothesis that drawing and mental models of 3D space are related is that people who were more accurate when drawing were also more accurate when judging the slant of edges in the still life. We were surprised to find, though, that participants were significantly *better* at slant judgement from the imagined indirect view. Slant is the pitch of an edge with respect to the horizontal, a value that is independent of viewpoint. The pitch of a roof doesn't change when the viewer moves. The slope of a road is the same whether viewed from the top or the bottom of the hill. Logically, the slant of an edge of an object in a still life is the same in both the direct and indirect conditions. We did not bring that fact to participants' attention and none of them mentioned it spontaneously. We cannot conclusively account for this finding. However, we chose edges for slant judgement arbitrarily after constructing the still lifes. That limited our options and so it may be that the edges judged in the direct and indirect conditions are not comparable. The accurate values of slant for the direct condition (52, 56 and 68 degrees) were all smaller than the accurate values for the indirect condition (71, 76 and 81 degrees).

We also found that experts were better than novices at projected angle judgement. This was a slightly unexpected finding. Although projected angle judgement appears to be analogous to authentic perspective drawing practice, the previous work (Carson & Allard, 2008, in press) that motivated this study had found no expert/novice difference. That experiment tested judgement of two edges; this tested six (three in each of the direct and indirect conditions). It may be that both of these novel perceptual judgement tasks (slant and projected angle) vary enough with angle size that it takes repeated measures and a well-chosen variety of edges to detect signal amidst the noise, which mirrors the rest of the angle judgement literature.

The design of this study could have been improved by counter-balancing the order of the drawing and perceptual judgement tasks and the order of the direct and indirect tasks. If the participants had been identified as expert or novice a priori, they could have been age-matched and the task order could have been counter-balanced within groups. If the participant pool had been carefully recruited for a balance on drawing expertise, sex, age and education in the STEM disciplines, we could have reduced the confounds in testing mental rotation.

In the absence of a gold-standard measure of expertise we were forced to rely on the conventional ten-year threshold of experience. Our drawing accuracy tasks exposes the weaknesses of using this threshold. Not all our experts were accurate and not all our novices were poor. Our participant pool included an excellent professional abstract painter, whose drawing accuracy was poor, and a nine-year “novice” who delivered the most accurate drawing of all (Appendix B.2). These cases highlight the advantage of using an objective performance-based measure when studying performance properties. We may improve the strength, clarity and specificity of findings by grouping participants on the basis of a more relevant dimension than experience.

Though angle accuracy was effective, it didn’t capture all the interesting properties of the drawings. For example, in the indirect condition, objects were often plausibly drawn but poorly positioned. Counting missing angles flagged this problem but did not describe or quantify it. In the next chapter, we will move on to a more complex geometric accuracy measure which solves this problem and does not limit us to analysis of drawings of mainly straight-edged subjects.

In a demanding, ecologically valid perspective drawing task, experts with ten or more years of drawing experience outperformed novices on an objective, angled-based measure of overall drawing accuracy. Having validated an angle-error approach, we demonstrated its utility for local analysis within the drawing. Accuracy varied with angle size—acute or obtuse—and provided evidence of a top-down influence of drawing context. Several novel perceptual tasks contributed converging evidence that drawing skill is linked to a superior understanding of the orientation and relationships in three-dimensional space of elements of the drawing subject.

Chapter 3

Geometric characterisation of drawings

3.1 Introduction

In the previous chapter, we developed an angle-based measure of drawing accuracy which was significantly associated with experience. This gave us a continuous measure of the overall accuracy of a drawing but also let us study *local* accuracy, such as the large divergent perspective errors people made when drawing the corner of a book. We showed that drawing accuracy can be usefully, objectively measured. The angle error approach does have limitations. It is well-suited to the measurement of accuracy of straight lines but not curved ones. That excludes most of portraiture, figure drawing and landscape imagery from analysis. It also seemed that, in practice, the relative positions of objects, their neighbours, the background and foreground were more flawed than drawings of the objects themselves. Angle error failed to capture or characterise those flaws. In the present work, we expand the level of analysis of accuracy from the marks in the drawing to the shapes they attempt to represent and discriminate between landmarks that have shifted (errors of occlusion) and those that are absent altogether.

To analyze the shapes in a drawing, we will consider an image as a collection of landmark points. Any element of that drawing—a teapot here, a tabletop there—can be represented by the landmarks it includes. If we consider those landmarks to be the vertices of a polygon, we can measure the geometric- properties of that polygon and compare it to the same element in a ground truth photograph of the original scene. Is the polygon in the drawing oriented at the same angle as its ground truth counterpart, shaped in the same proportions, scaled appropriately, and drawn in the same position? Representing perspective drawings as deformations of the geometry of a ground truth photograph of the original scene will allow us to assess drawing accuracy quantitatively, without the judgment of human critics. The advance here would not be in eliminating qualitative assessment, but in using the analysis of geometric deformations to study the accuracy of drawings at higher levels of organization. We can make comparisons not just between lines but between properties of the scene and performances of the participants. We can, for example, begin to study what sorts of objects are harder or easier to draw accurately, and, from that, what properties of human visual perception might make this so.

In this study, we applied objective geometric measures of drawing accuracy to the drawings of thirty-four participants who attempted to draw a three-dimensional still life. Those participants varied widely in experience. By visual inspection, the drawings appeared to range in quality from novice to expert. We controlled the scene being drawn and the viewpoint from which participants drew it. Both of these are ecological constraints that are consistent with practice in the domain. Artists must control their viewpoint to draw accurately from observation (Nicolaides, 1941), and the still life is a staple of drawing instruction and practice (Edwards, 1979).

Our goal was to develop these measures as a tool for behavioural psychological research into drawing and visual perception. In such research it would be valuable to characterise drawing error globally (across the drawing as a whole) and locally (in regions of particular research interest). We hypothesized that the global error measurements of the drawings would discriminate correctly between expert and novice participants. Comparing the performance of experts and novices is an economical way to identify the skills and behaviour that are necessary for excellence in a domain (Ericsson et al., 1993). Our previous studies of perspective drawings (Carson & Allard, 2008; Carson, Quehl, Aliev, & Danckert, 2011) found that angles were drawn more accurately by expert participants than novices, a finding that suggests objective measures of geometric deformations are consistent with expertise levels. However, an angle-based approach is of somewhat limited utility, best suited to the analysis of drawings of predominantly straight-edged subjects. Here we expanded that approach to include polygon-based measures that can be extended to drawings of naturalistic subjects. For a measure of global drawing accuracy to be useful, it should not only discriminate between experts and novices, but provide a finer-grained, continuous measure of the drawing. For the present study, we compared an overall measure of drawing accuracy to the number of years of drawing experience of participants rather than just to two groups—experts and novices.

The same studies of perspective drawings (Carson & Allard, 2008; Carson et al., 2011) found that some angles were drawn significantly more accurately than others. This suggests that not all objects are equally difficult to draw. We hypothesized that local error measurements within the drawings would confirm that error is not uniform across the image and would suggest some features or properties of the image are harder to draw than others. We represented the drawings as sets of interest points in the image and selected groups of interest points to study as polygonal Regions of Interest (ROIs). For a measure of local drawing accuracy to be useful, it should be sensitive enough to characterise differences in accuracy between well-chosen ROIs. For the present study, we chose ROIs to explore a fundamental question in the study of drawing: Are people more accurate when drawing objects (the "positive space") than the spaces between them (the "negative space")? That is, are people better at drawing things than at arranging those things correctly in the image?

3.2 Method

3.2.1 Model

To quantify drawing accuracy, we need to be able to measure the deformation of drawings in a data set by comparison to a reference photograph representing the ground truth of the drawing subject. We also want to be able to measure the deformation of ROIs within the drawings and to begin to characterize the ways in which they are deformed. We want to describe (and measure) how ROIs are oriented, proportioned, scaled, and positioned in a drawing, but we do not propose to exhaustively measure the deformation of all possible ROIs in the image. Instead, the model allows the researchers to choose ROIs of specific relevance to their own analyses.

We can represent the ground truth photograph and each of the drawings as collections of interest points corresponding to unambiguous landmarks throughout the image. We can then tally the number of missing interest points—a simple measure of completion. Next we can tally the number of places where one edge should occlude or intersect another to create an interest point, but fails to.

We choose N polygonal ROIs, ($Q_{1 \ldots N}$ for $N \geq 3$; Figure 3.1), whose vertices are interest points on the ground truth photograph. Each Q_j has area A_j, k sides of length $L_{j,k}$ and k rays of length $R_{j,k}$ from its centroid C_j to the k^{th} vertex. Each Q_j has orientation α_j, which is the direction (in radians) of the minor principal axis of the second moment of the area of Q_j. We chose the minor rather than the major principal axis because it aligns graphically with the largest dimension of the ROI but also—since the designation of major and minor principle axes is unstable for objects that are close to square—to choose the principal axis that minimizes orientation error, $\mathcal{O}_{i,j}$ (defined below). This model accurately reports the orientation error as long as the orientation error is less than 90 degrees. This is a reasonable constraint to place on the suitability of drawings for analysis and was true for all the ROIs in all the drawings examined in this study.

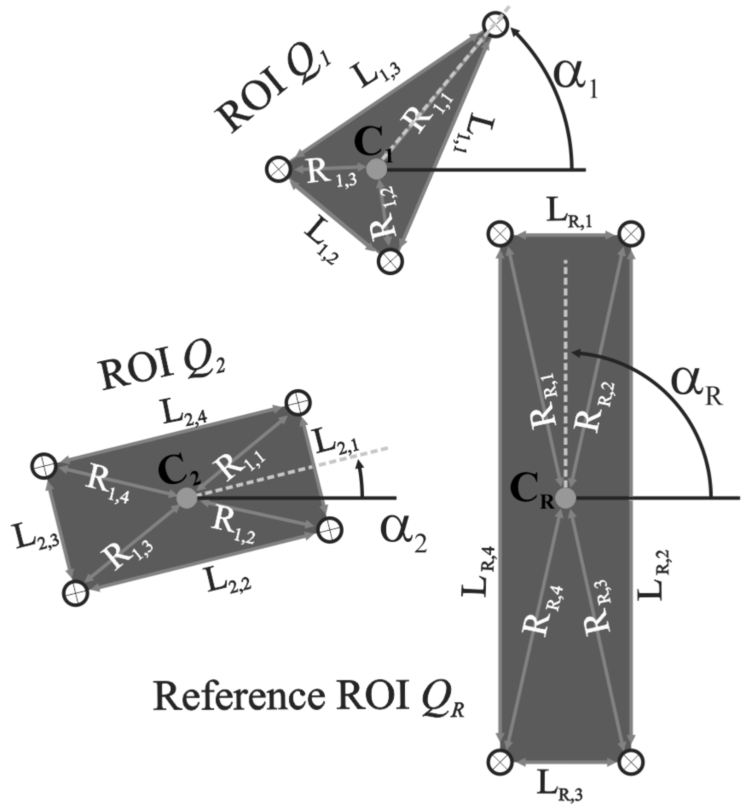

Figure 3.1: In the model, the j Regions Of Interest (ROIs) in each drawing are described by j polygons (Q_j), one of which is designated as the reference ROI (polygon Q_R). The orientation (α), area (A), centroid (C), side lengths (L), and the ray lengths (distance from the centroid to each corner, R) are used to compare the proportionality, scaling, orientation and positions of the polygons in the i^{th} drawing to their counterparts in a ground truth photograph. Polygon properties such as $L_{j,k}$ represent the k^{th} L of the j^{th} polygonal ROI.

By inspection of the data set, we choose a useful reference ROI, Q_R, which is present and well proportioned in every drawing, has a moderately large area A_R, and a long narrow shape that will produce stable principal axes. Informally, the reference ROI Q_R should be a large long shape that was demonstrably easy for all the participants to draw accurately. For every drawing, we measure four properties of each ROI Q_j (where that ROI is present in the drawing):

1. its orientation error relative to the reference ROI Q_R;
2. its proportionality error *independent* of the reference ROI;
3. its scaling error relative to the reference ROI Q_R; and,
4. its position error relative to the reference ROI Q_R.

The intuitive meanings of these four error measures are described graphically in Appendix A as a complement to the formulas here. Finally, we calculate an overall measure of the polygon-wise accuracy of the drawing, which is the mean of the four types of polygon errors across all the ROIs in the drawing, each normalized to the largest error of its kind across the i drawings.

Our primary interest is in ROI analysis but we also want to validate the present results against previous work with these drawings (Carson & Allard, 2008). From the model, it is simple to compare angles of interest between drawings and the ground truth photograph. An angle β is formed by three non-collinear interest points and we measure its error by comparison to the same angle in the ground truth photograph.

Polygon orientation error

The orientation error $\mathcal{O}_{i,j}$ (see also, Appendix A) quantifies how well the j^{th} ROI, Q_j, is oriented in the i^{th} drawing in comparison to the ground truth photograph. Since participants drew at different orientations relative to the edge of the paper, the orientation of each j^{th} ROI is calculated relative to the orientation of the reference ROI (α_R in Figure 3.1). (Starred variables describe the ground truth photograph.)

$$\mathcal{O}_{i,j} = \left|(\alpha_{i,j} - \alpha_{i,R}) - \left(\alpha^*_j - \alpha^*_R\right)\right| \quad (3.1)$$

$$\mathbb{O}_i = \sum_j \mathcal{O}_{i,j} \quad (3.2)$$

$\mathbb{O}_i$ is the total orientation error, in radians, across all the ROIs, $Q_{1...N}$, in the i^{th} drawing. Since the orientation error is relative to the reference ROI, $\mathbb{O}_i$ does not include an error measure for Q_R. $\mathbb{O}_{i,j}$ will have a value from zero (when all of the ROIs in the i^{th} drawing are perfectly oriented relative to the reference ROI) to $\frac{\pi}{2}(N-1)$ (when all of the ROIs in the i^{th} drawing are misoriented by 90 degrees relative to the reference ROI, an unlikely extreme; the model assumes that orientation error is less than 90 degrees for all ROIs).

Polygon proportionality error

The proportionality error $\mathcal{P}_{i,j}$ (see also, Appendix A) is a scale- and orientation-invariant measure that quantifies how accurately the shape of the j^{th} ROI, Q_j, is represented in the i^{th} drawing in comparison to the ground truth photograph, independent of all other ROIs in the drawing. The proportionality error of Q_j in the i^{th} drawing is calculated by taking the sum of the differences between the k^{th} normalized perimeter segment (Figure 3.1; $L_{i,j,k}/\sum_w L_{i,j,w}$) and the corresponding perimeter segment on the ground truth photograph ($L^*_{j,k}/\sum_w L^*_{j,w}$), and summing it with the differences between the k^{th} normalized ray length ($R_{i,j,k}/\sum_w R_{i,j,w}$) and the corresponding ray length on the ground truth photograph ($R^*_{j,k}/\sum_w R^*_{j,w}$).

The perimeter and ray segments are both needed in this measure in order to capture errors of stretch and shear. The sums are halved to make the maximum error of this measure exactly one.

$$\begin{aligned}\mathcal{P}_{i,j} &= \sum_k \frac{1}{2}\left|\frac{L_{i,j,k}}{\sum_w L_{i,j,w}} - \frac{L^*_{j,k}}{\sum_w L^*_{j,w}}\right| \\ &\quad + \sum_k \frac{1}{2}\left|\frac{R_{i,j,k}}{\sum_w R_{i,j,w}} - \frac{R^*_{j,k}}{\sum_w R^*_{j,w}}\right| \end{aligned} \tag{3.3}$$

$$\mathbb{P}_i = \sum_j \mathcal{P}_{i,j} \tag{3.4}$$

$\mathbb{P}_i$ is the total proportionality error across all the ROIs, $Q_{1\ldots N}$, in the i^{th} drawing. Since the proportionality error is scale- and orientation-invariant, $\mathbb{P}_i$ also includes an error measure for the reference ROI, Q_R. $\mathbb{P}_i$ will have a value from zero (i.e., all of the ROIs in the i^{th} drawing, including the reference ROI, are perfectly proportioned) to N (i.e., all of the ROIs in the i^{th} drawing, including the reference ROI, are maximally misshapen).

Polygon scaling error

The scaling error $\mathcal{S}_{i,j}$ (see also, Appendix A) quantifies how accurate the size of the j^{th} ROI, Q_j, is in the i^{th} drawing in comparison to the ground truth photograph. Since participants drew at different scales, the size of the drawing is normalized with respect to the root of the area of the reference ROI (Q_R in Figure 3.1), which would be the height of a square of equal area. This reference length is intentionally dependent on the area of the reference polygon—rather than its height—to make normalization less sensitive to deformations of the shape of the reference polygon. The scaling error is calculated by taking the magnitude of the difference between the normalized area A of j^{th} ROI in the i^{th} drawing, $A_{i,j}/A_{i,R}$, and its corresponding ROI in the ground truth photograph, A^*_j/A^*_R, and dividing that by the area of the corresponding ROI, A^*_j/A^*_R, to yield a measure of the percentage area error.

$$\mathcal{S}_{i,j} = \frac{\left|\frac{A_{i,j}}{A_{i,R}} - \frac{A^*_j}{A^*_R}\right|}{\frac{A^*_j}{A^*_R}} \tag{3.5}$$

$$\mathbb{S}_i = \sum_j \mathcal{S}_{i,j} \tag{3.6}$$

$\mathbb{S}_i$ is the total scaling error across all the ROIs, $Q_{1\ldots N}$, in the i^{th} drawing. Since the scaling error is relative to the reference ROI, $\mathbb{S}_i$ does not include an error measure for Q_R. $\mathbb{S}_i$ will have a value from zero (i.e., all of the ROIs in the i^{th} drawing are perfectly scaled relative to the reference ROI) to a large positive value of no fixed upper bound (i.e., all of the ROIs in the i^{th} drawing are extremely mis-scaled—too large or too small—relative to the reference ROI).

Polygon position error

The position error $\mathcal{T}_{i,j}$ (see also, Appendix A) quantifies how well the j^{th} ROI, Q_j, is positioned in the i^{th} drawing relative to the reference ROI, Q_R. Since participants drew at different scales, the size of the drawing is normalized with respect to the root of the area of the reference ROI (Q_R in Figure 3.1). As in scaling error, this reference length is used to make normalization less sensitive to deformations of the shape of the reference polygon. Since participants drew at different orientations relative to the edge of the paper, the orientation of each j^{th} ROI is calculated relative to the orientation of the reference ROI (α_R in Figure 3.1). Finally, since the choice of reference ROI is well-informed but arbitrary, and a satisfactory position error should be as independent of that choice as possible, the position error is also normalized with respect to the ground truth distance between the centroids of the j^{th} ROI and the reference ROI. The position error of the (normalized, re-oriented) j^{th} ROI is the distance from its centroid, C_j, to the centroid, C_j^*, of its corresponding ROI in the ground truth photograph, divided by the distance between the centroids, C_j^*, and C_R^*.

$$\gamma_i = \alpha_{i,R} - \alpha_R^* \tag{3.7}$$

$$\mathcal{T}_{i,j} = \frac{\left\| \begin{bmatrix} \cos\gamma_i & \sin\gamma_i \\ -\sin\gamma_i & \cos\gamma_i \end{bmatrix} \frac{\vec{C}_{i,j} - \vec{C}_{i,R}}{\sqrt{A_{i,R}}} - \frac{\vec{C}_j^* - \vec{C}_R^*}{\sqrt{A_R^*}} \right\|_2}{\left\| \frac{\vec{C}_j^* - \vec{C}_R^*}{\sqrt{A_R^*}} \right\|_2} \tag{3.8}$$

$$\mathbb{T}_i = \sum_j \mathcal{T}_{i,j} \tag{3.9}$$

$\mathbb{T}_i$ is the total position error across all the ROIs, $Q_{1...N}$, in the i^{th} drawing. Since the position error is relative to the reference ROI, $\mathbb{T}_i$ does not include an error measure for Q_R. $\mathbb{T}_i$ will have a value from zero (i.e., all of the ROIs in the i^{th} drawing are perfectly positioned relative to the reference ROI) to a large positive value of no fixed upper bound (i.e., all of the ROIs in the i^{th} drawing are extremely mis-positioned relative to the reference ROI).

Angle error

Separately from polygon analysis, we may analyze angles of interest. We measure the magnitude of the error, in radians, of every n^{th} angle of interest $\beta_{i,n}$ in the i^{th} drawing against the corresponding angle in the ground truth photograph β_n^*.

$$\mathcal{B}_{i,n} = |\beta_{i,n} - \beta_n^*| \tag{3.10}$$

$$\mathbb{B}_i = \sum_n \mathcal{B}_{i,n} \tag{3.11}$$

$\mathbb{B}_i$ is the total angle error across all the angles of interest in the i^{th} drawing. $\mathbb{B}_i$ will have a value from zero (i.e., all of the angles of interest in the i^{th} drawing are perfectly accurate) to a large positive value bounded by $2\pi n$ (i.e., all of angles of interest in the i^{th} drawing are extremely inaccurate).

Overall polygon error

Finally, an overall measure of polygon accuracy ($\mathbb{G}$) is calculated for each i^{th} drawing against the other drawings. For this analysis, we assume—for lack of any evidence or hypothesis to the contrary—that all the error dimensions are equally important. We define the overall error of the i^{th} drawing as the mean of its four polygon errors (orientation $\mathbb{O}_i$, proportionality $\mathbb{P}_i$, scaling $\mathbb{S}_i$, and position $\mathbb{T}_i$), each normalized to the largest error of its kind across all i drawings to make the maximum error of this measure exactly one.

$$\mathbb{G}_i = \frac{1}{4}\left[\frac{\mathbb{O}_i}{max\,\mathbb{O}_i} + \frac{\mathbb{P}_i}{max\,\mathbb{P}_i} + \frac{\mathbb{S}_i}{max\,\mathbb{S}_i} + \frac{\mathbb{T}_i}{max\,\mathbb{T}_i}\right] \tag{3.12}$$

$\mathbb{G}_i$ is the overall polygon error in the i^{th} drawing. $\mathbb{G}_i$ will have a value from zero (i.e., all of the ROIs in the i^{th} drawing are perfectly oriented, proportioned, scaled and positioned) to one (i.e., each of the ROIs in the i^{th} drawing is the worst of its kind, on all four dimensions of error, across all the drawings in the dataset).

3.2.2 Participants

For a previous study (Carson & Allard, 2008) 34 adult participants (11 men, 23 women; age M= 38.97, SD= 18.8) were recruited to make perspective drawings from a still life. We revisited their drawings for this chapter.

Many of the participants were recruited from the university student population and some earned course credit for their participation in the original study. We also recruited participants from the local art community in order to find drawing experts. We used a survey of the participants' drawing experience to identify ten experts (seven men, three women; i.e., people with at least ten years experience).

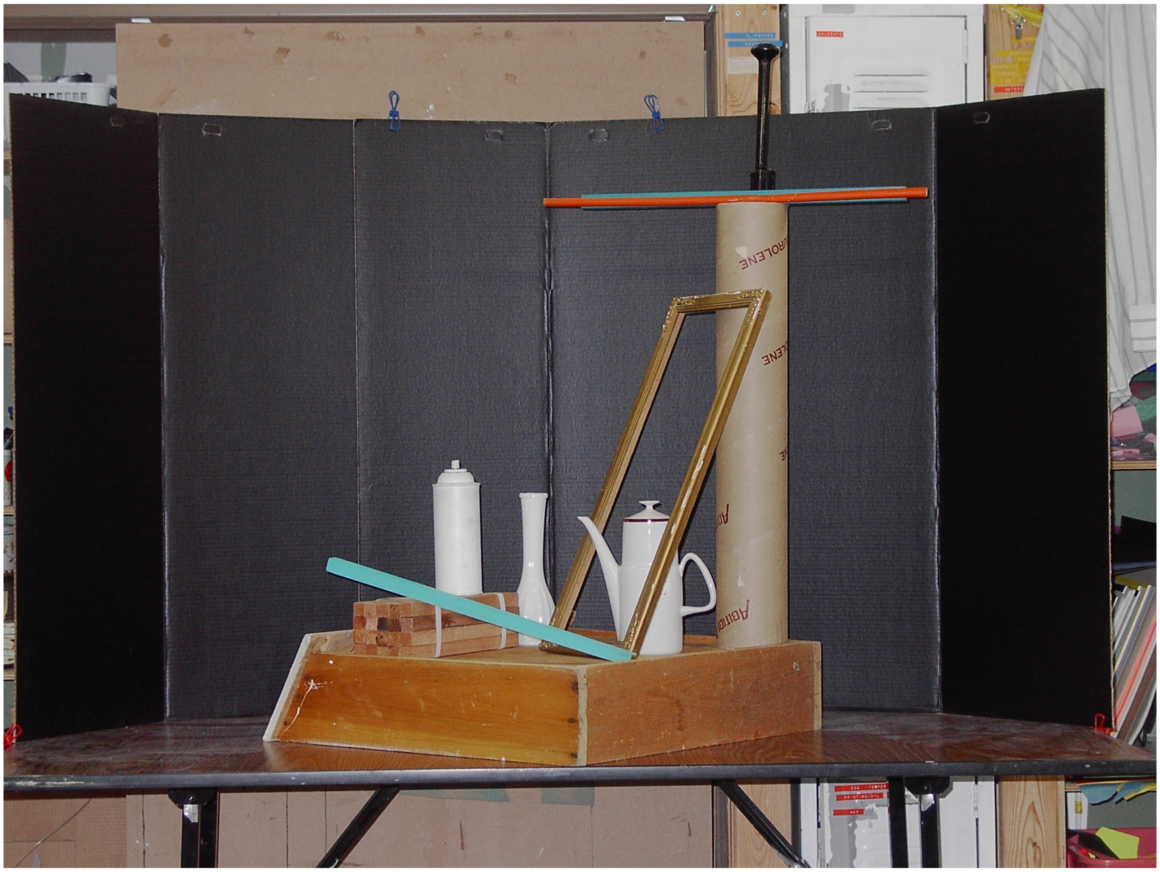

Figure 3.2: The still life Alpha, direct view, i.e., as seen from the participant's viewpoint

3.2.3 Apparatus and Procedure

This was an ecologically valid drawing task, designed to be as close to studio practice of still life drawing as was practical. The subject of the drawings was a complex three-dimensional still life (Figure 3.2), which was thoroughly but unobtrusively marked so we could reconstruct it reliably for each participant. The drawings were made in a variety of locations under available general-purpose lighting. We conclude that the illumination and contrast were adequate because none of the participants complained about the lighting, or appeared to have any difficulty in seeing the still life and all its detail clearly enough to draw. The still life was challenging to draw—including many objects, overlapping objects, and objects receding from the viewer in space—in order to be certain that even the expert artists would produce some measurable errors. We conclude from the subsequent analysis, and from the number of good-natured complaints during the task about the complexity of the subject, that we were successful in this aim.

The still life included a small visual landmark we used to coach participants on how to locate and maintain the desired viewpoint consistently (Figure 2.2). The participants drew, without time limit, from direct observation of the still life while seated roughly three metres away. They used 2B pencil (and eraser) on an $8\frac{1}{2}$ by 11 inch piece of commonplace 20-pound white paper, sometimes supported by a clipboard and sometimes resting directly on a table.

We took a ground truth photograph from the same viewpoint as that from which all the participants drew (Figure 3.2). We identified 96 points of interest on the ground truth photograph, still life landmarks such as the corners of objects and occlusion points where the edge of an object in the still life met or overlapped the edge of an object behind it. We chose unambiguous landmarks so that we could clearly identify the corresponding points in a drawing, where they existed. Landmarks were manually annotated on the photograph using a Matlab script to establish the ground truth for our analysis. These annotations were repeated and compared with the same annotations by a second rater to ensure that the error between annotations was small. This is consistent with the practice of previous researchers doing manual annotation of still images (Winter, 2005). Drawings were manually annotated using a Matlab script to capture the interest points across the drawings wherever they existed. To annotate the drawings, we superimposed an image of the ground truth data as vertices and edges over a scanned drawing then dragged the vertices to their deformed locations in the drawing.

Even drawing experts are not perfect. There were missing interest points in most drawings, each of which was tagged as an error of omission. Occlusion points in the still life, represented in the image by the intersection of two edges, were assigned two interest points, one associated with each of the edges. In a drawing where those edges intersected incorrectly (or failed to intersect at all), those interest points were tagged as having an occlusion error. We analyzed errors of omission and occlusion, and omitted incomplete ROIs from polygon-wise analysis. This annotation gives us a collection of interest points representing the (deformed) locations of landmarks in each drawing.

We chose as the reference ROI Q_5, a rectangle corresponding to a tall cylinder in the still life. By inspection, the cylinder was present in all drawings and was drawn relatively accurately by all participants. We chose four non-overlapping ROIs (Figure 3.3). The ROI Q_1 corresponds to the left front face of the base of the still life; Q_2 corresponds to a spray paint can on the left side of the still life; Q_3 describes the small space between the spray can and a small vase; and Q_4 describes the large open space above the picture frame and to the left of the cylinder. We chose just one angle of interest, β_1, the foreground corner of the base, to validate the present results against previous work with these drawings (Carson & Allard, 2008). For each of the drawings, we computed all of the errors described in the model: polygon-wise errors of orientation, proportionality, scaling and position; overall polygon error; and, the angle error of β_1.

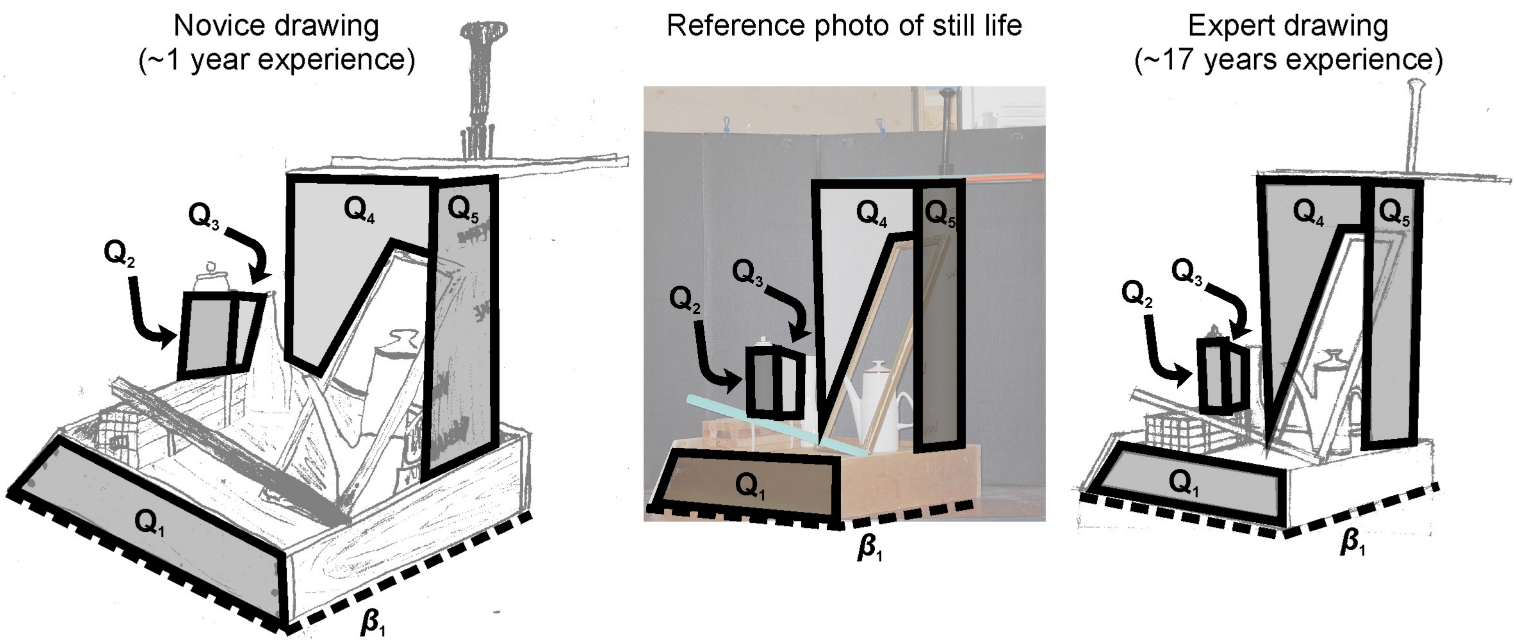

Figure 3.3: The ground truth photograph (centre) is bracketed by examples of a novice drawing (left: one year of drawing experience; overall polygon error $\mathbb{G} = .70$) and an expert drawing (right: 17 years of drawing experience; overall polygon error $\mathbb{G} = .16$), each marked with their Regions of Interest ($Q_{1...4}$), the reference ROI (Q_5), and the angle of interest (β_1). For this illustration, all three images have been approximately resized to show the reference ROI, Q_5, as the same height.

3.3 Results

3.3.1 Overall polygon errors

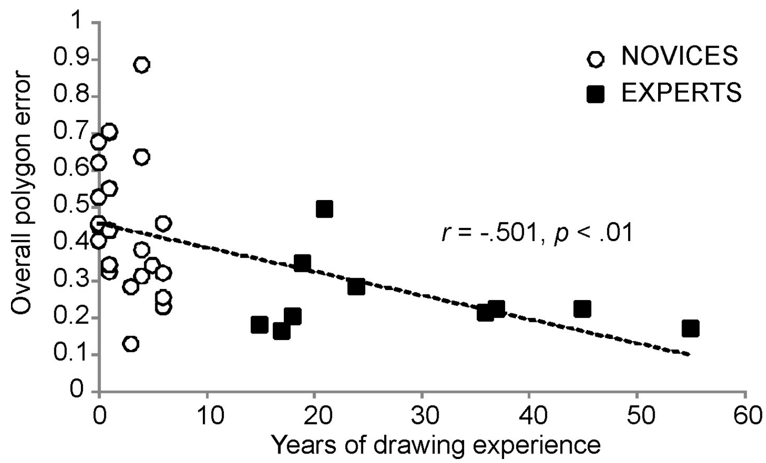

Figure 3.4: Overall polygon error measure, $\mathbb{G}_i$, as a function of a participant's years of drawing training and experience.

Experts had significantly lower overall polygon error (M= 0.249, SD= .101) than novices (M= 0.447, SD= .181; $t(32) = 3.229$, $p<$ .01). The overall polygon error, $\mathbb{G}_i$, correlated strongly with a participant's years of drawing experience and training, $r(34) = -.501$, $p<$ .01 (Figure 3.4). The more experience participants had, the lower their overall polygon error. There was also a strong correlation between experience and error when drawing angle β_1, $r(34) = -.440$, $p<$ .01 (Figure 3.5). The more experienced participants drew this angle more accurately, which is consistent with the previous analysis of these drawings (Carson & Allard, 2008).

3.3.2 Errors of omission and occlusion

Overall, experts committed significantly fewer errors of occlusion (M= 9.4, SD= 4.43) than novices (M= 15.5, SD= 5.32; $F(1, 32) = 10.17$, $p<$ .01, η^2= .241). Experts omitted fewer interest points (M= 1.3, SD= 1.34) than novices (M= 2.3, SD= 1.74) but this difference was not significant (Figure 3.6).

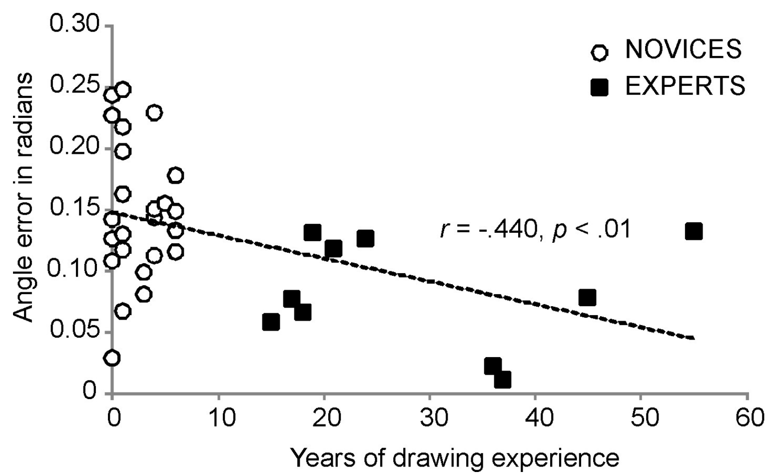

Figure 3.5: Error drawing angle β_1 as a function of a participant's years of drawing training and experience.

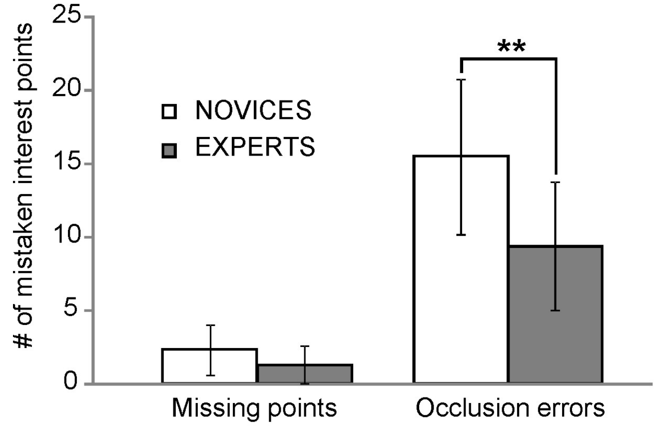

Figure 3.6: Number of missing points (left) and number of occlusion errors (right) as a function of participants' level of experience. Error bars represent ± one standard deviation.

3.3.3 Polygon-wise errors

Across all the participants, expert and novice, the polygon-wise errors of the drawings were closely related to one another (Figure 3.7).

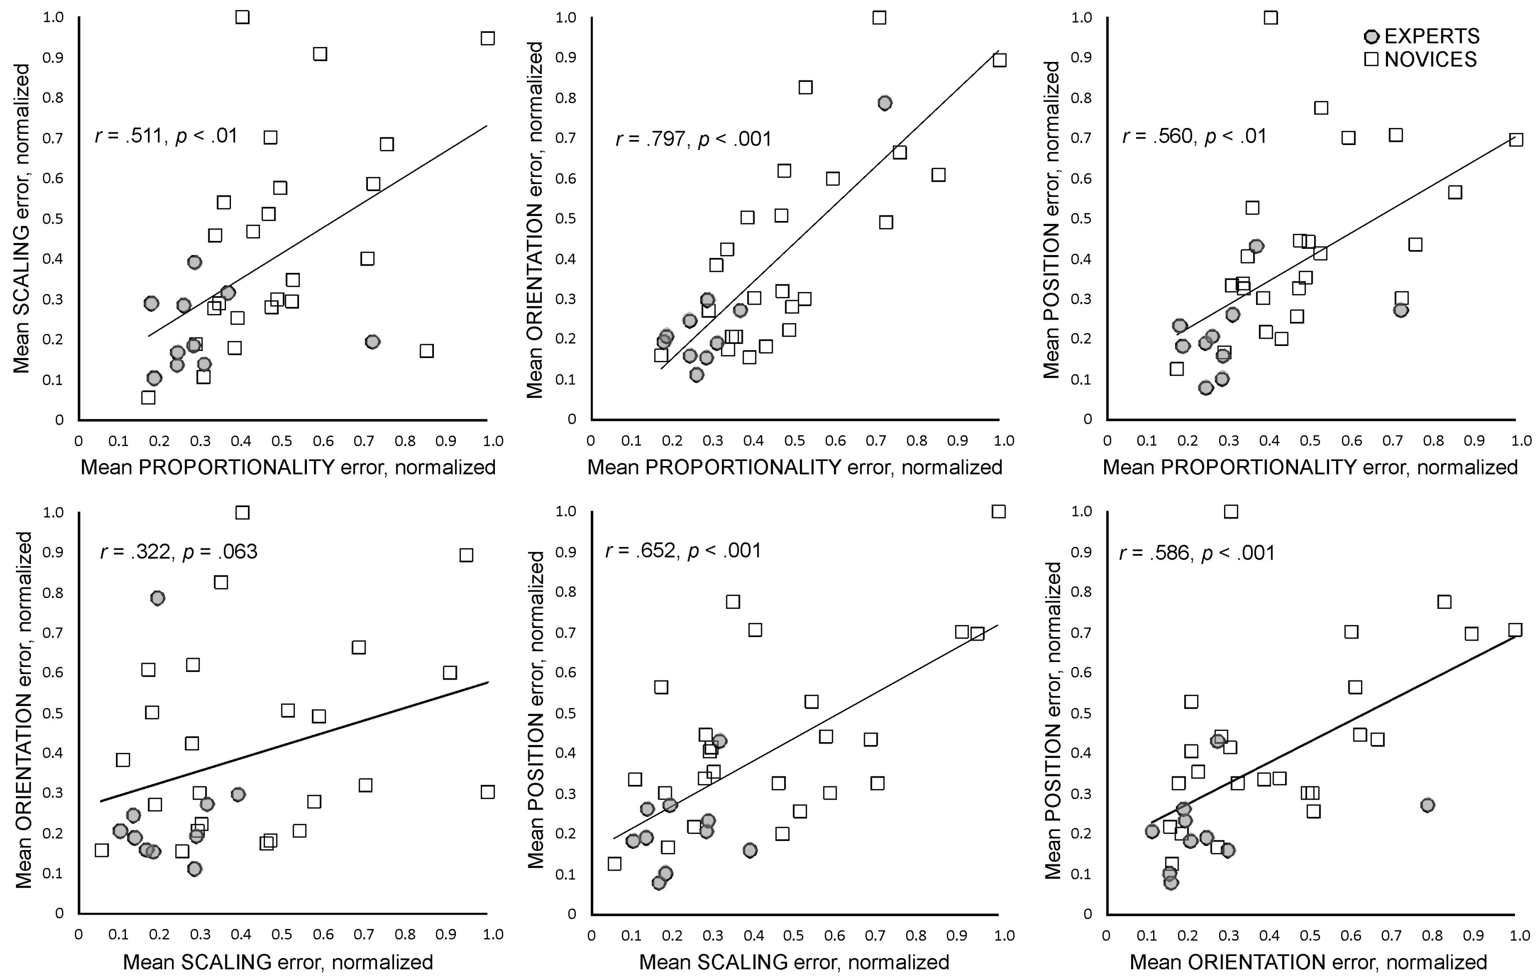

Figure 3.7: The polygon-wise drawing errors of proportionality, scaling, orientation and position were closely related to one another.

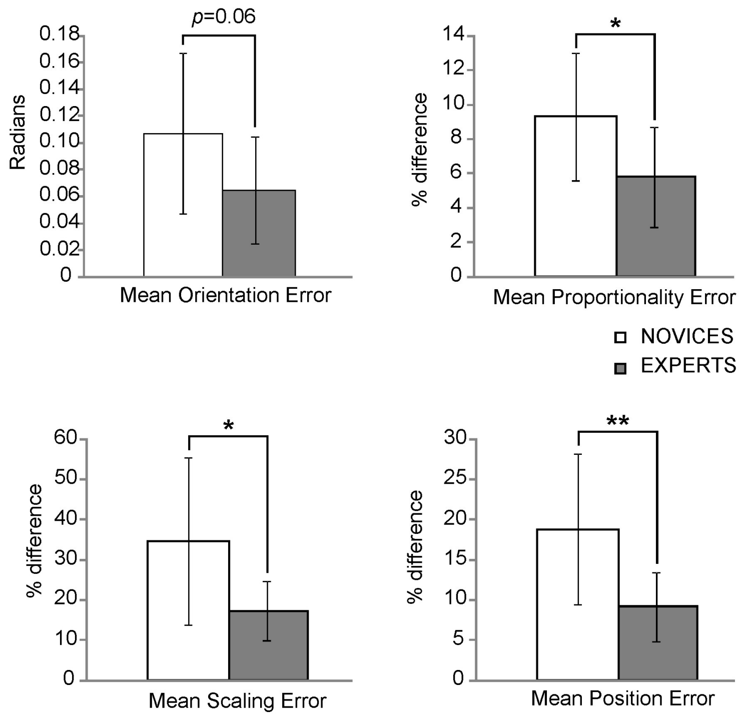

Figure 3.8: Experts made smaller polygon-wise errors than novices. Error bars represent $\pm$ one standard deviation.

Experts were significantly more accurate than novices at capturing the proportions of the polygons in their drawings ($F(1, 32) = 6.99$, $p < .05$, $\eta^2 = .179$; Figure 3.8), at scaling polygons ($F(1, 32) = 6.46$, $p < .05$, $\eta^2 = .168$) and at positioning the polygons ($F(1, 32) = 9.63$, $p < .01$, $\eta^2 = .231$). Experts made smaller errors of orientation ($M = 0.06$, $SD = 0.048$) than the novices ($M = .10$, $SD = 0.06$), an advantage that approached significance ($F(1, 32) = 3.74$, $p = .062$, $\eta^2 = .105$).

3.3.4 Validating the choice of reference polygon

In order to confirm that the analysis was robust, we repeated the error measurements using Q_1 as the reference ROI. The results, reported in Appendix D, confirmed that the choice of a reference ROI had minimal effect on the findings.

3.3.5 Positive and negative space

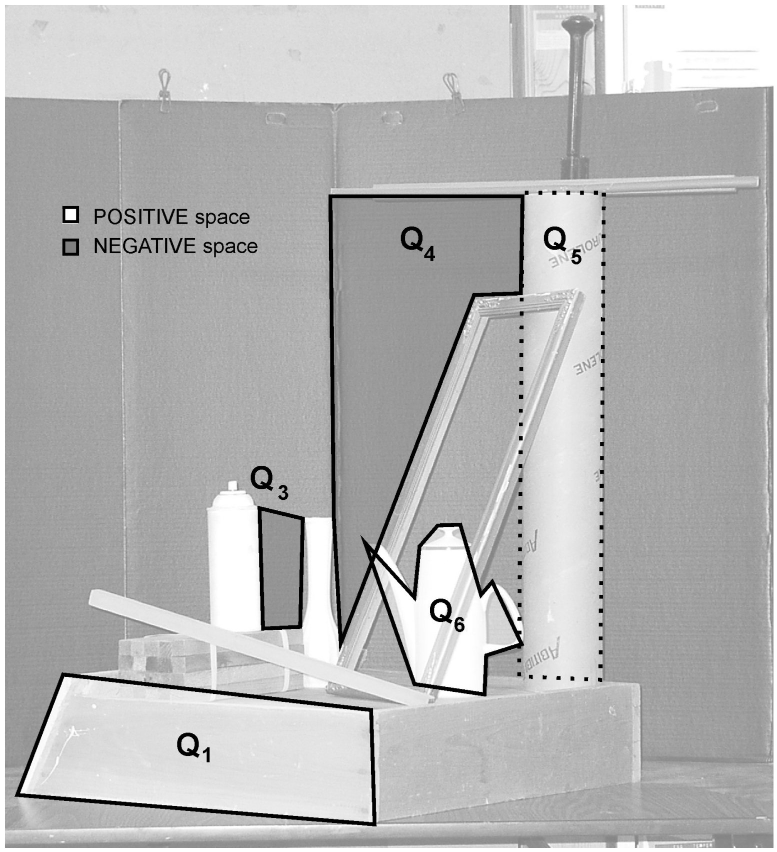

Figure 3.9: Polygons Q_1 and Q_6 describe objects, or positive space. Polygons Q_3 and Q_4 describe regions between objects, or negative space.

One ROI was added to compare drawing performance within participants on polygons in positive and negative space: Q_6 captures the main landmarks of the teapot, a positive space polygon, and is more complex than the polygon it replaces, Q_2. The proportions of positive space polygons were drawn significantly more accurately than negative space polygons by all participants ($F(1, 31) = 10.33$, $p< .01$, $\eta^2= .250$). The positions of positive space polygons were drawn more accurately, a difference that approached significance ($F(1, 31) = 3.99$, $p= .055$, $\eta^2= .114$). The orientation of positive space polygons was significantly *less accurate* than that of negative space polygons ($F(1, 31) = 7.39$, $p< .05$, $\eta^2= .193$; Figure 3.10). There was no significant difference between positive and negative space in the accuracy of polygon scale (Figure 3.10). Furthermore, there was no significant interaction

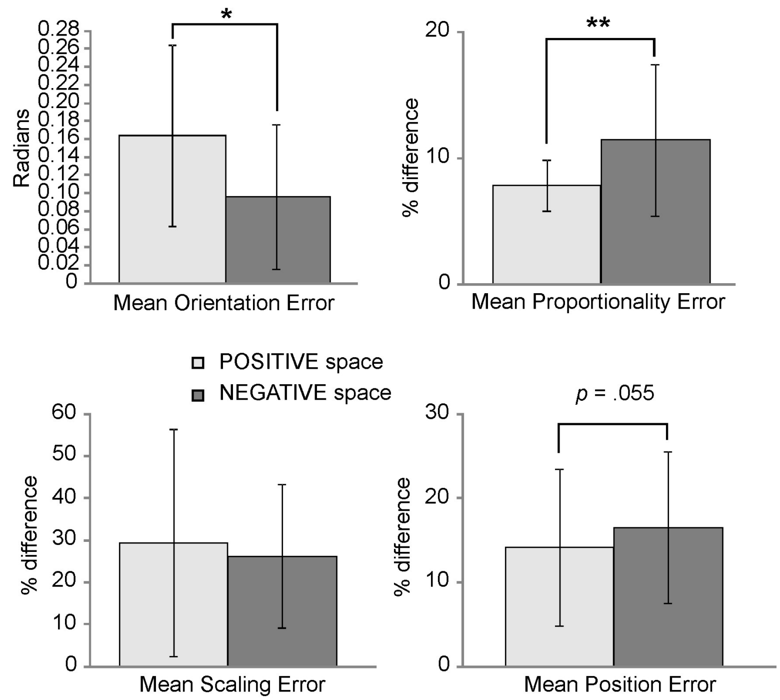

Figure 3.10: Proportionality and position error were higher for negative space polygons than for positive space polygons. Orientation error was *lower* in the negative space. Error bars represent ± one standard deviation.

between these errors and expertise.

3.4 Discussion

No one will be surprised that a veteran artist draws better than an inexperienced student. The contribution of this work is not in recognizing drawing excellence but in measuring it in new, objective ways that support varied and complex analyses. We have proposed and validated reasonable geometric measures of drawing accuracy that are practical, useful, and consistent with the apparent excellence of the drawings and the reported expertise of the participants (Figures 3.3, 3.4, and 3.5). These measures are based not on piece-by-piece mark-making but on geometric properties of shapes in the image, an approach that is more ecologically valid than previous work and allows us to study shapes chosen for their utility to particular research questions.

Figure 3.11: Lunia Czechowska, 1919, by Amedeo Modigliani (left) and Mystery and Melancholy of a Street, 1914, by Giorgio de Chirico (right).

In this geometric model of drawing accuracy, any element of an image can be represented by a set of landmarks and the polygon that is formed when those landmarks are connected. The correctness of that polygon in a drawing can be well described by comparing it to the ground truth of the original subject with four accuracy measures. We can summarize those measures informally: orientation reports on how much the shape leans or tilts; proportionality reports on the correctness of the shape and how badly it may have been squashed, stretched or dented; scale reports on how much the shape may have been shrunken or inflated relative to its neighbours; and, position reports on how close the shape is to being in the right place in the drawing. With these measures we can be specific and objective about the ways in which a drawing distorts its subject. In the visual arts, some distortions are deliberate and expressive, such as the distinctive elongation of Modigliani's portraits—which would be characterized by its proportionality measure—and the conflicting vanishing points of de Chirico's arcades—which would be described by the orientation measure (Figure 3.11). Our model allows us to tease apart these different dimensions of drawing accuracy.

The overall drawing accuracy correlated with years of drawing training and experience, showing that measured accuracy of a perspective drawing captures expertise. Furthermore, experts outperformed novices in every dimension, which assures us that each of the four polygon-specific measures also captures expertise. It remains to validate these measures against qualitative assessments of the same drawings to determine how much of a drawing's "goodness" is accounted for by its accuracy.

We can also see that, while experience correlates with accuracy, it does not guarantee it. Novices certainly varied widely in their performance. However, a few experts were quite inaccurate, and the most accurate performance came from a participant with just three years of drawing training. Looking at the drawing with the lowest error (Appendix B.1) strengthens our confidence in our accuracy measures. Despite the participant's *relative* inexperience, the drawing is both measurably and qualitatively accurate. The variability in the performance of experts and of novices is not altogether surprising given the inherent fallibility of self-report measures of experience. Here we tallied years of experience generously and uncritically, making it likely that expertise would sometimes be misstated. For example, we did not distinguish between, or verify, years of experience as a hobbyist and years of professional illustration. We tallied years of experience in areas of art practice (such as sculpture and abstract painting) that might or might not include deliberate practice in drawing. Finally, we had no reliable way to encourage a participant to report years of deliberate practice outside the classroom and the workplace. These were conscious, conservative choices to stack the deck in favour of the null hypothesis, but as a consequence we expected to see some designated experts under-perform and some designated novices do well. This only highlights the original motivation for the study—to develop a performance-based measure of drawing excellence as a basis for analysis.

One of the goals of this research is to invent measures of drawing accuracy as tools for the development and validation of drawing exercises. Drawing error is not uniform across the image. That is, some parts of an image are harder to draw than others. For example, drawing teachers have long observed that students depict foreground objects—the positive space—more accurately than the background and the spaces between objects—the negative space (Edwards, 1979; Nicolaides, 1941). Here we selected ROIs to investigate this claim by comparing polygon accuracy in the positive and negative space in the drawings of all our participants. Results confirm the conventional wisdom of drawing teachers. Negative space polygons are more poorly proportioned and positioned than positive space polygons for all participants, even experts. Interestingly, negative space accuracy is not worse on all dimensions. People are equally accurate—or inaccurate—in the positive and negative space when scaling shapes, and they were *more* accurate orienting shapes in the negative space. This may be a clue to which properties of a representation matter most to viewers. We do not tolerate errors of proportion or position in the positive space that we overlook in the negative space, but we are relatively insensitive to scale between the negative and positive space. This is consistent with the literature on the perception of size constancy (e.g., Gregory, 1963). From this literature we would predict that viewers are unlikely to notice errors of scale of objects, and in our analysis, people were no more sensitive to scale errors of objects (the positive space) than of the negative space.

The psychologist and the artist read these data differently. To the psychologist, they show that, if our drawings represent what we notice about what we see, then the important dimensions of the figure over ground are proportionality and position. To the artist, on the other hand, these data are signposts to where we can improve our drawing; as long suspected, we need to pay more attention to the shapes of the negative spaces. Having confirmed that people draw less accurately in the negative space and that we can measurably detect the difference, we propose, in future work, to use this geometric model to test the effectiveness of negative space drawing exercises to see which, if any, measurably improves performance. The orientation finding merits more investigation to determine whether it is, perhaps, a trade-off for proportional and position accuracy in the positive space.

An important strength of this geometric model is that it should generalize to naturalistic subjects, which will support more varied and ecologically valid research into drawings of curvilinear and biomorphic forms such as landscape, still life and the figure. This is essential for applied research in drawing education. In future work, we plan to confirm this by extending this study to naturalistic subjects.

With this approach to analysis in hand, it will be possible to construct controlled still lifes—building in properties of particular research interest *a priori*—and to measure the effects of changes on local and global drawing accuracy.

It will also be possible to construct drawings with controlled amounts of specific types of error to investigate visual perception of perspective images. We found, in the first study in Chapter 2, significant evidence of divergent perspective when even the most experienced and accurate participants drew the book (Section 2.3.3). Cubes and other rectangular prisms seem to be especially prone to distortion. In our geometric model, such a distorted rectangular prism—represented as a polygonal ROI—would have a fairly accurate position and a slightly skewed orientation. However, its size would be too large and the proportionality would be wrong (Figure 3.12). People presumably believe those drawings are satisfactorily accurate. How wrong would such an image have to be—standing on its own, or in direct comparison to a ground truth photograph—for the average person to detect the error?

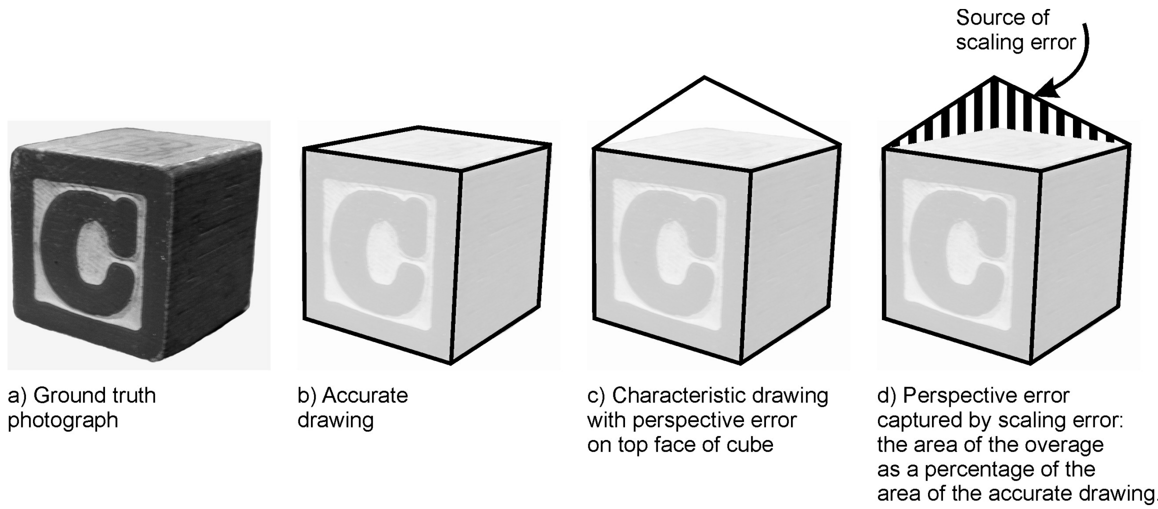

Figure 3.12: People characteristically distort the top face of the cube which—in our geometric model—would be captured by errors of proportionality and scaling.

Drawing is a delightfully complex visuomotor task and there are many likely sources of error: in perception of the subject, in the planning and execution of the drawing, in the evaluation of that drawing and in attempts to correct its shortcomings. Furthermore, drawing is a serial task where early errors will be compounded. When we can quantify and localize drawing error, we can begin to attribute that error appropriately to perceptual processes and drawing skills, which will tell us more about how we see and draw the world.

Chapter 4

Qualitative ratings of drawings

4.1 Introduction

In the previous chapters, we developed two different measures of drawing accuracy and validated them against the expertise of the drawers. Accuracy proved to be a fruitful basis for analysis of the properties of the drawings and of the skills of the participants. In the present study, we set out to demonstrate a relationship between the accuracy and the *quality* of drawings. We predicted that more accurate drawings would be rated of higher quality. Our accuracy measures will be more credible and useful when we can show that they are consistent with the literature and the ratings of human critics.

The history of research into drawing has been almost entirely a history of ratings by human critics. Researchers have asked critics to rate different properties of drawings or drawing performance, such as quality in a general sense (McManus et al., 2010), accuracy (Cohen & Earls, 2010), or, rarely, specific aesthetic properties (Kozbelt & Serafin, 2009). Most researchers have recruited raters with some specialist knowledge of drawing, such as art professors (Ostrofsky, Kozbelt, & Seidel, 2012) and practising artists (Csikszentmihalyi & Getzels, 1971). It takes some effort to find suitable experts and it may not always be necessary. Cohen and Jones (2008) found that a large pool of novices produced comparable accuracy ratings of portrait drawings to those of experts. However, novice and expert raters are not always interchangeable. When Kozbelt and Serafin (2009) asked novices and experts to make quality ratings of drawings in progress, they found that novices tended to equate quality with accuracy while experts equated quality with originality.

We are interested in objective measures of drawing accuracy to *complement* quality ratings, not to replace them. There are research questions for which ratings from human critics are problematic. For example, critics usually rate the drawing as a whole, which limits us to analysis of global problems. In Chapter 2, local accuracy measures pinpointed a problem with divergent perspective and in Chapter 3, local accuracy measures confirmed conventional studio wisdom about positive and negative space, two findings that would not be possible with global ratings of drawings. Human critics are authoritative on aesthetic judgements, but they are as susceptible to perceptual biases as the drawers so they may contribute confounds as well as insight to questions of the role of visual perception in drawing. The ratings of human critics are ordinal and coarse-grained—the most detailed scale in the literature is only a 12-point scale, used by McManus and colleagues (2011)—which limits analysis.

In this study, we invited people to grade the quality of the drawings obtained through the research in the previous chapters, but offered no further direction or guidance about what that might encompass. Participants were not blind to the goal of the study, as we told them "We've already measured the accuracy of these drawings and we're trying to figure out how much accuracy has to do with overall quality of the drawings." We did not define quality as beauty or any aesthetic property. We deliberately left the participants to interpret quality according to their own preferences. We showed participants a photograph of each still life before they began grading, which presumably creates some inclination in the viewer to assess the drawings on the basis of how successfully they represent the still life. However, the research question of interest was not, *Are human raters effective at judging drawing accuracy?* but the more foundational question, *Are accurate perspective drawings seen as being any "good"?*

4.2 Method

4.2.1 Participants

We recruited 106 adult participants (47 men, 58 women and one who declined to report sex; age $M = 20.64$, $SD = 5.4$) to judge perspective drawings. Most of the participants were university students earning course credit for their participation. Some participants were recruited on Twitter, from the local art community and through word of mouth. Three participants reported that they drew in some professional capacity. None taught drawing. Eight participants reported ten or more years of training and experience with drawing. None of the raters had participated in the study from which the drawings were collected.

4.2.2 Apparatus

The study was conducted on-line. Participants were told that they were to rate the overall quality of drawings that had already been graded for accuracy. They rated 104 drawings (34 drawings of still life Alpha, randomly ordered, Figure 4.1; 35 randomly-ordered drawings of still life Bravo; and, 35 randomly-ordered drawings of still life Charlie;) and then did a survey.

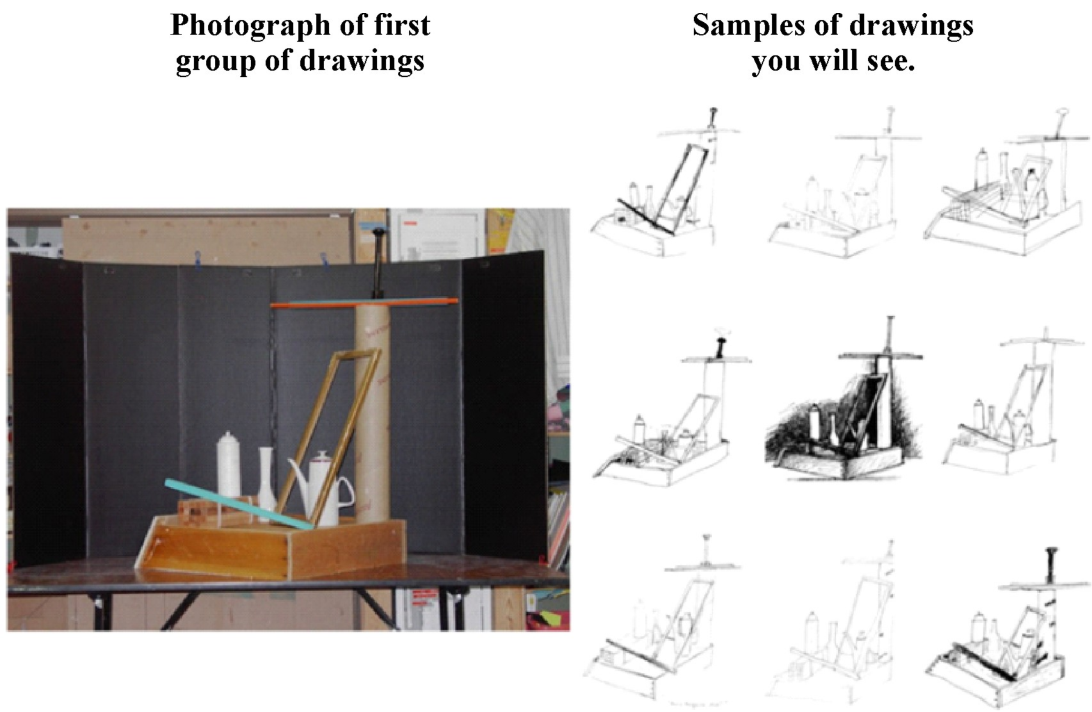

Figure 4.1: The ground truth photograph of the still life Alpha and a varied selection of sample drawings.

4.2.3 Procedure

Using the web survey, participants graded the quality of 104 drawings on a scale from 1 ("Poor") to 10 ("Excellent"). After giving informed consent, participants performed several tasks. Participants worked solo and it usually took about an hour for each of them to complete all of the study tasks.

We introduced the rating task and alerted the participants that there were three sets of drawings to rate, each based on a different still life. They graded the drawings of still life Alpha, then Bravo, then Charlie. We describe the procedure for Alpha. The other two were identical.

We showed participants a photograph of the still life Alpha, alongside nine sample drawings from the dataset (Figure 4.1). Those drawings were chosen in advance to represent a wide range of performance. This gave participants a chance to reflect, before they started grading, on what they might assess as "poor" or "excellent." Then we showed the 34 drawings, one at a time, in random order, and collected participants' ratings.

The drawings of Alpha were completed from direct observation with unlimited time, the drawings of Bravo were also completed from direct observation, but under a ten-minute time constraint, and the drawings of Charlie were completed in ten minutes, from an imagined, indirect (side) view. Consequently, we predicted that the drawings of Alpha would be rated the most highly overall and the drawings of Charlie would be lowest.

After the participants had graded the drawings in all three sets, we collected demographic information, including information about their drawing training and experience.

4.3 Results

4.3.1 Reliability of quality ratings

A reliability analysis showed that the participants' ratings were remarkably consistent with one another (Cronbach's $\alpha = .98$). 11 participants declined to rate one or more drawings, for a total of 261 (2.3%) missing ratings out of 11,024 possible ratings.

4.3.2 Ratings by drawing novices and experts

We asked raters to report how much experience they themselves had with drawing. We identified raters as experts ($N = 8$) if they reported ten or more years of training and drawing experience. When we compared the ratings made by experts and novices overall, we found that novices ($N = 98$) rated drawings more positively (M= 4.34, SD= 0.96) than the experts did (M= 3.50, SD= 1.22), $F(1, 104) = 5.364$, $p< .05$, $\eta^2 = .048$. The novice ratings, though higher, agreed with the expert ratings across the drawings, $r = .932$, $p< .001$ (Figure 4.2). Having concluded that the drawing experience of the rater had little effect on the ratings, we collapsed the two groups of raters (expert and novice) for subsequent analyses.

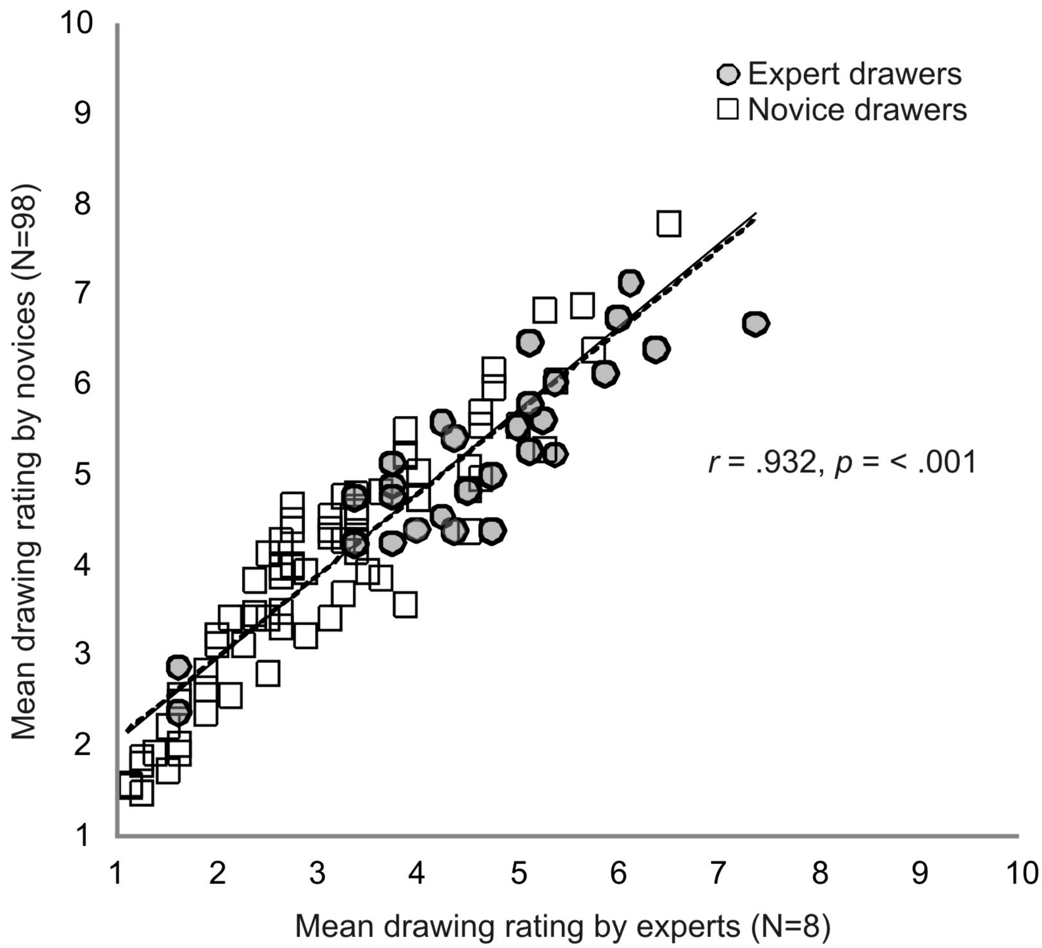

Figure 4.2: Drawing quality as rated by novices as a function of drawing quality as rated by experts.

4.3.3 Quality across the three drawing sets

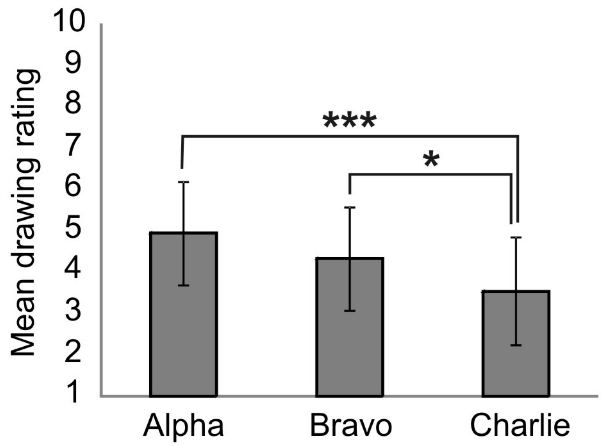

Figure 4.3: Quality rating as a function of the specific still life task. Error bars represent ± one standard deviation.

For each drawing, we calculated a mean rating of quality across all the raters ($N = 106$) and used that as a dependent variable.

As a check of the usefulness of the quality ratings, we compared the mean of those drawing ratings across each of the three sets of drawings. As we had predicted, the drawings of Alpha were rated most highly (M= 4.96, SD= 1.25), Bravo slightly lower (M= 4.32, SD= 1.25), and Charlie lowest of all (M= 3.54, SD= 1.30). The difference between datasets was significant ($F(2, 101) = 10.820$, $p<$ .001, η^2= .176) and after Bonferroni correction, results showed that Charlie differed significantly from Alpha ($p<$ .001) and Bravo ($p<$ .05) but Alpha and Bravo did not differ significantly from each other (Figure 4.3).

4.3.4 Drawer's experience and drawing quality

Across all of the drawings of Alpha, Bravo and Charlie, the experience of the person who made the drawing was a significant predictor of its quality rating, $r(104) = .311$, $p<$ .001 (Figure 4.4). The more experience the drawer had, the more highly rated the drawing.

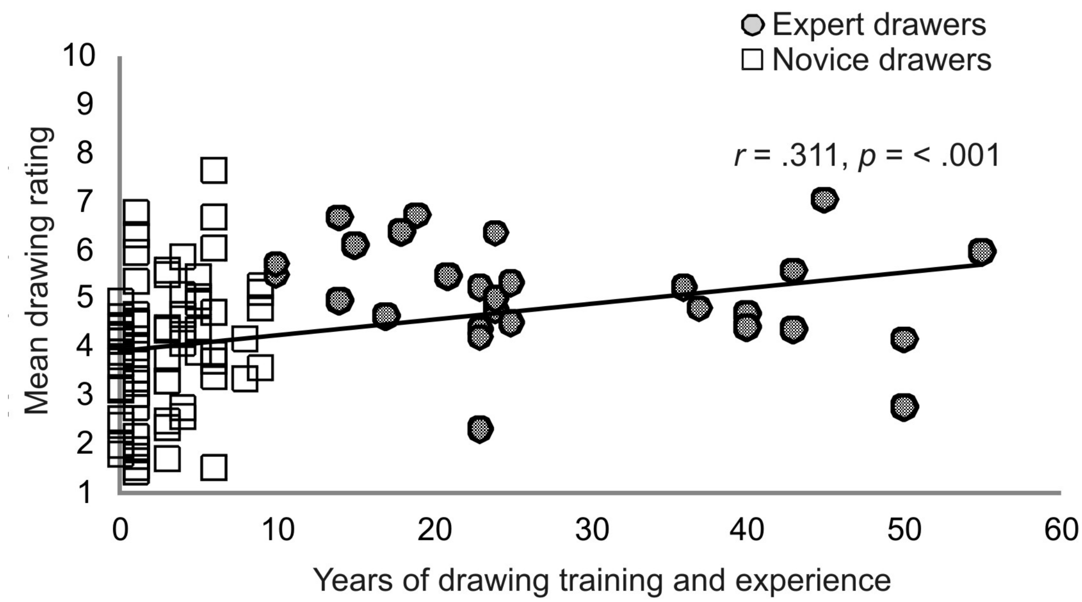

Figure 4.4: Quality rating as a function of the years of drawing training and experience of the drawer for the drawings of still lifes Alpha, Bravo and Charlie.

4.3.5 Angle accuracy and drawing quality

In the analysis of Chapter 2, we calculated an angle-based accuracy measure for the drawings of still life Bravo and still life Charlie: the mean of the percentage magnitude angle drawing error. Considering the drawings as a single set of 70 (on the grounds that their accuracy was measured in the same way), we compared the drawings' accuracy and quality ratings. The mean of the percentage magnitude angle drawing error across a drawing and its mean quality rating by all the raters correlated significantly ($r(35) = -.224$, $p < .05$, one-tailed), such that drawings rated to be of higher quality had lower mean percentage magnitude angle drawing error (Figure 4.5).

4.3.6 Polygon accuracy and drawing quality

In the analysis of Chapter 3, we calculated a polygon-based accuracy measure for the 34 drawings of still life Alpha. For each of the drawings, we compared its overall polygon error and its mean quality rating by all the raters.

They correlated at a level that approached significance, $r(34) = -.273$, $p = .059$, one-

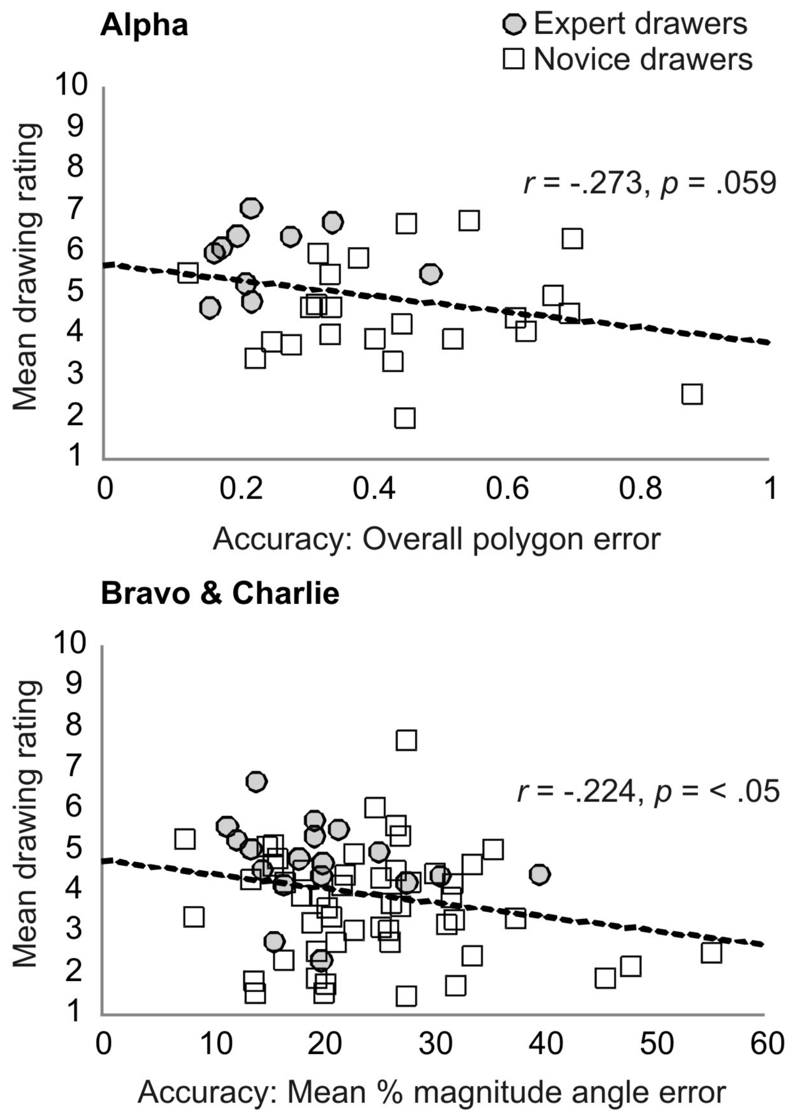

Figure 4.5: Quality rating as a function of polygon error for the 34 drawings of still life Alpha (upper) and as a function of angle error for the 70 drawings of still lifes Bravo and Charlie (below).

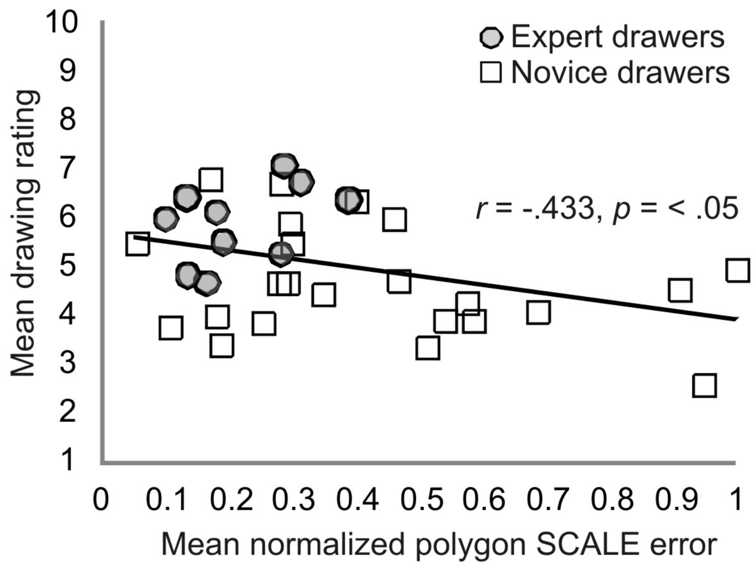

Figure 4.6: Quality rating as a function of mean polygon scale error (here, normalized).

tailed, such that drawings rated to be of higher quality had lower overall polygon error. This analysis may not have reached significance because it was based on a smaller sample size than the analysis (above) of drawings measured with angle-based accuracy ($N = 70$).

We also compared the mean quality of drawings with their four mean polygon-wise errors. Quality did not correlate with orientation, proportionality or position error, but mean quality rating correlated significantly with the mean polygon scale error, $r(34) = -.433$, $p < .05$ (Figure 4.6). Drawings with more accurately scaled polygons were rated of higher quality.

4.4 Discussion

As we predicted, the quality of drawings was consistent with the measured accuracy of those drawings (Figure 4.7). Raters generally agreed with one another on how "good" drawings were. The people who had spent a lot of time drawing agreed, too, but set a slightly more demanding standard for quality. Not only were more accurate drawings better-rated, but the drawings made by people with more experience were generally better-rated than those of novices. That supports the consistency of our accuracy measures with both experience and quality ratings.

However, if we compare the highest-rated drawings to the most accurate, it is clear there is more to quality than just accuracy. For example, the most accurate drawing of still life Alpha ($G = .12$; Figure 4.7, upper right) had a crisp charm and was well-rated at 5.50 (on a scale from 1, or "Poor," to 10, or "Excellent"), but still stood only 12^{th} in quality in the pack of 34 drawings. Its most highly rated rival (7.05; Figure 4.7, upper left) was quite accurate, $G = .22$, but only 7^{th} in accuracy overall.

It should also be noted that, while accuracy and experience are consistent with one another *overall*, they are not interchangeable. One of the main motivations for these studies was to find and validate objective measures of drawing accuracy that might be more reliable yardsticks of expertise than retrospective self-report of drawing experience. The most accurate drawing in Alpha (Figure 4.7, upper right) was made by a participant with just three years of reported drawing experience. The reader can clearly see that this was an accurate drawing (see also Appendix B.1, where the drawing is overlaid on the ground truth photograph). The least accurate (Figure 4.7, lower right), was made by a participant with *four* years of reported drawing experience. The reader can tell at a glance that the drawing is inaccurate. Based on the participants' reports of their years of drawing training and experience, a classic expert/novice measure would fail to discriminate these two drawings, but polygon accuracy places them leagues apart. The quality ratings of human critics—who rated the most accurate drawing at 5.50 and the least accurate drawing at 2.58—suggest that accuracy could be a better yardstick of expertise than experience.

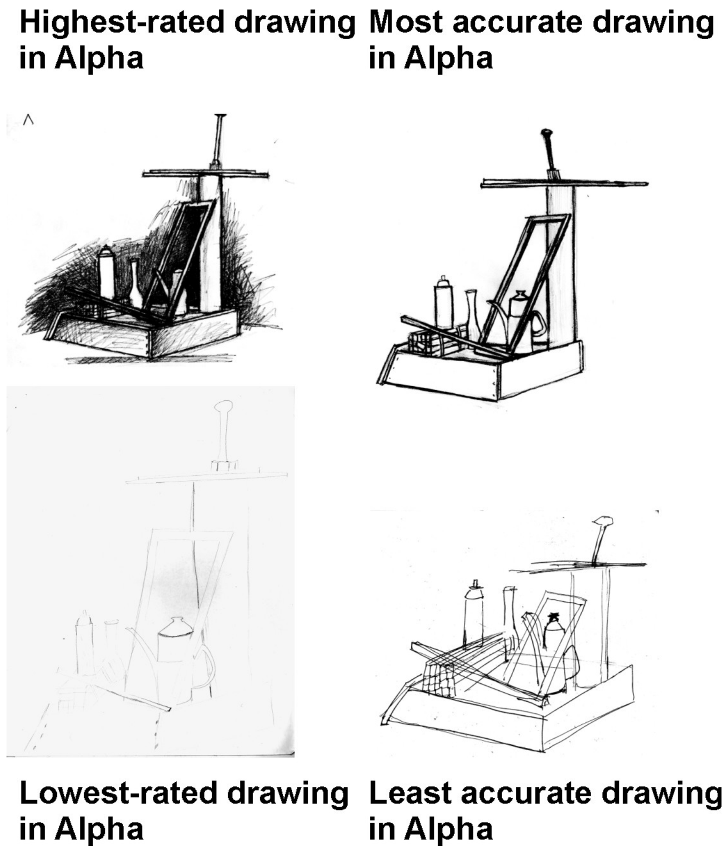

Figure 4.7: Upper left panel, the most highly-rated drawing of still life Alpha; upper right panel, the most accurate drawing; lower left panel, the lowest-rated drawing; and, lower right, the least accurate drawing

Chapter 5

Neural correlates of drawing in an artist with neglect

All of the findings reported here are behavioural but when we see differences in visuomotor skills that appear to dissociate the *look* of objects from their *location*, we must wonder about possible neural correlates. In 1983, Mishkin, Ungerleider, and Macko described the existence in the brain of two extrastriate pathways processing visual information in parallel. The ventral stream, they proposed, was specialized for "object vision" and the dorsal stream was specialized for "spatial vision." These became known, respectively, as the "what" and "where" pathways. Goodale and Milner (1992) suggested that the distinction was more functional, that the ventral stream processed vision for perception and the dorsal stream processed vision for action, and, informally, these became the "what" and "how" pathways. Vision, though seamlessly integrated in everyday experience, does more than one thing at a time with the same stimuli. It appears to support object identification in the ventral stream and visually-guided motor behaviour in the dorsal stream. In this context, drawing may be a very special behaviour. Drawing is a finely-coordinated visuomotor activity that both relies on and interprets visual stimuli—real, remembered and imagined—then reverse engineers the percepts to represent them as marks on the page. Furthermore, imagining a still life from another viewpoint requires mentally shifting from one frame of reference to another. This thesis provided some evidence that artists are better at this shift, supporting a better mental model of 3D space, and perhaps an important role for the dorsal stream in drawing.

The neural correlates of drawing have been the subject of recent research interest (Bhattacharya & Petsche, 2005; Chatterjee, 2006; Kottlow, Praeg, Luethy, & Jancke,

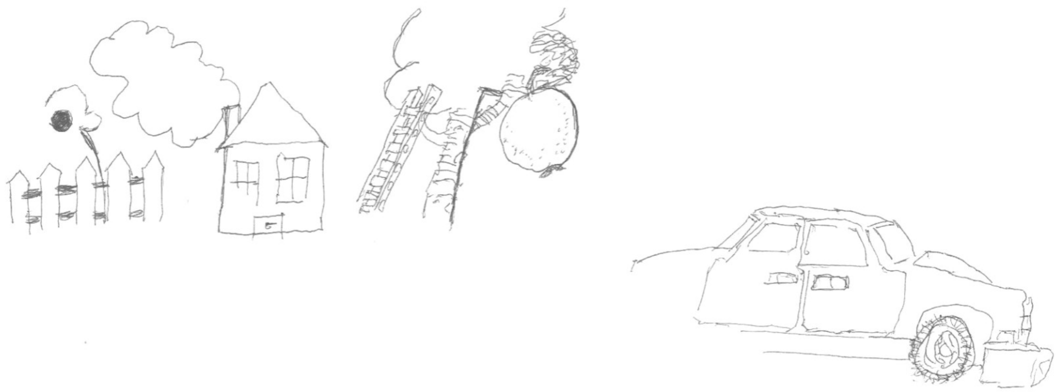

Figure 5.1: Figure copying task by DM, three months after a right hemisphere stroke.

2011; Kozbelt & Seeley, 2007; Solso, 2001) though the research into dorsal stream involvement has been primarily focused on constructional apraxia rather than artistic practice (Makuuchi, Kaminaga, & Sugishita, 2003). Of particular value may be recent work on changing frames of reference for spatial cognition in healthy individuals (Committeri et al., 2004) and in people with hemispatial neglect (Yue, Song, Huo, & Wang, 2012) which suggests that different frames of reference are encoded by the dorsal stream alone (egocentric or viewer-centred) or an integration of dorsal and ventral streams (allocentric, whether object- or landmark-centred). Researchers have also divided three-dimensional space *behaviourally* into peripersonal space—near space within and on which we can physically act—and extrapersonal space—space beyond our reach which can be subdivided into focal, action, and ambient domains (Previc, 1998). There is evidence from patients with right hemisphere lesions of neglect in peripersonal but not extrapersonal space (Halligan & Marshall, 1991; Halligan, Fink, Marshall, & Vallar, 2003).

It could be that experienced artists develop better dorsal stream functioning, or better communication between the dorsal and ventral streams, and there is reason to believe this would be right hemisphere dominant. That is *not* to suggest that drawing is all "on the right side of the brain" as the popular DIY manual (Edwards, 1979) suggests. However, the neural correlates of the 3D model of space, which seems to be critical to perspective drawing, are predominantly right hemispheric (Adair & Barrett, 2008). That being so, we would expect damage there to interfere with expert drawing performance. We note that people with right hemisphere damage may present with hemispatial neglect, including making drawings from which they characteristically omit information on the left side. This is a serious drawing deficit regardless of expertise, but we further speculate that right parietal damage—particularly when it involves the dorsal regions which are associated with

Figure 5.2: Freehand butterfly drawing by DM, three months after a right hemisphere stroke.

peripersonal neglect of near space—would undermine an experienced artist's previously superior mental model of 3D space.

Three years before this study, a former illustrator and animator, DM, suffered a right hemisphere stroke and presented, three months after his stroke, with acute left unilateral neglect (shape cancellation 67.74% left omissions, 6.45% right omissions; mean line bisection error 2.16%; Figures 5.1 and 5.2). CT scans demonstrated that DM had suffered a right middle cerebral artery stroke affecting mainly subcortical regions, including the coronal radiata and the anterior limb of the internal capsule, extending posteriorly to the right external and extreme capsules. Damage also affected the lateral aspect of the right basal ganglia posteriorly and the medial aspect of the right temporal lobe. At the time of data collection for this study, DM, a right-handed 62-year-old man with vision corrected to normal and at least ten years of drawing experience, had not recovered use of his left side and showed slight signs of neglect visible only in figure copying (Figure 5.3) and qualitatively in how he performed the tasks. He was clearly aware of his history of neglect. As he completed the clock drawing, he critiqued his work in progress and, while doing a star cancellation task, he coached himself aloud to remember to proceed systematically and to look to his left. In the figure copying (Figure 5.3, left), the left point of the star and the left leaf on the daisy are smaller. On the clock face, the numbers on the right are crowded and the 6 was repeated.

If drawing practice strengthens an artist's model of 3D space, we would expect that a patient like DM—with some impairment of the dorsal stream or at the very least a disconnection between the dorsal and ventral streams in his lesion—would show some loss

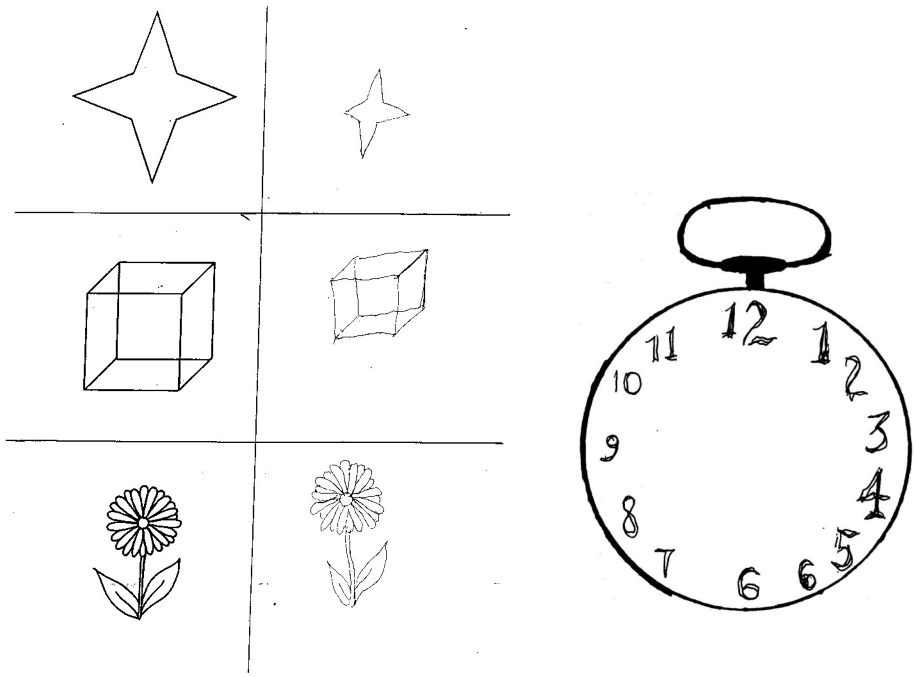

Figure 5.3: Figure copying task (left) and clock face drawing (right) by DM, three years after a right hemisphere stroke.

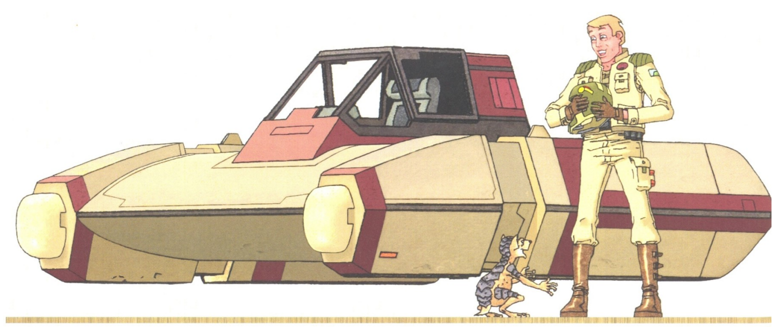

Figure 5.4: Drawing made by DM before his stroke.

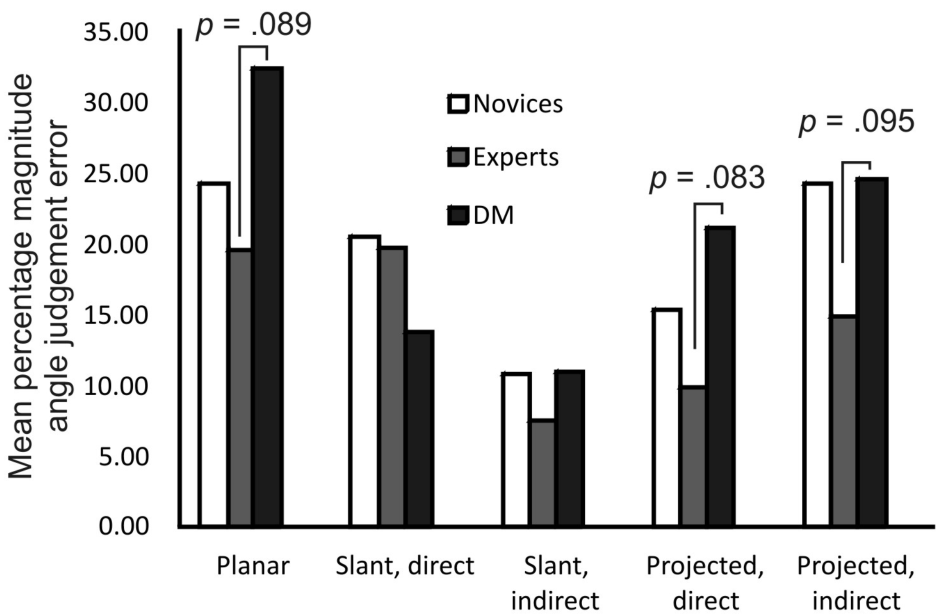

Figure 5.5: DM's performance on the angle judgement tasks compared with that of healthy experts and novices.

of detail on the left characteristic of his neglect but would also present with distortions of 3D space that would not have been apparent pre-injury (Figure 5.4).

DM agreed to participate in the drawing study with enthusiasm. Because he was still confined to bed, had no use of his left hand, and tired easily, we did not ask him to perform the mental rotation test or standard still life drawing tasks. With some modifications to accommodate the setting and his disability, DM performed the angle judgement and spatial extent tasks of the study in Chapter 2 and we compared his performance to healthy experts and novices. We set up the still lifes on a table over the foot of the bed and took new reference photographs from DM's viewpoint and from a side view to confirm the answers to the perceptual judgement tasks in this setting. When DM made perceptual judgements of slant, projected angles, extent and occlusion, he was much closer to the still life than other participants. He was, in fact, within arm's reach. He sometimes reached out with his right hand and touched the objects he was talking about, even across his midline.

Consistent with our prediction, and despite his historical drawing expertise, DM's angle judgement was less accurate than the expert and novice participants on almost all tasks (Figure 5.5). In order to compare DM's angle judgement performance with healthy ex-

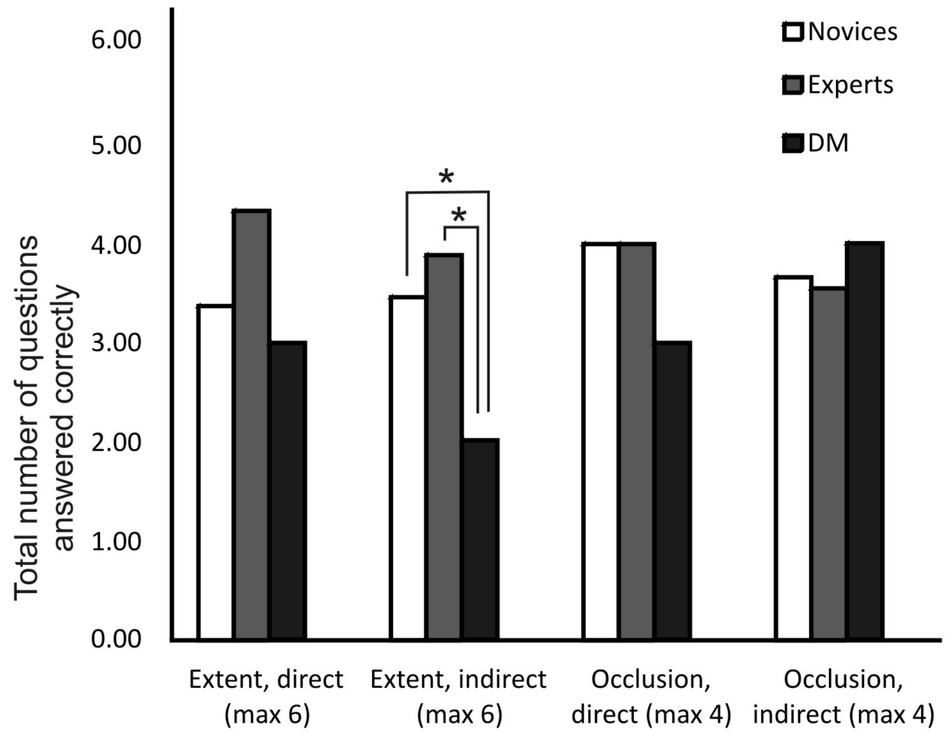

Figure 5.6: DM's performance on the extent and occlusion judgement tasks compared with that of healthy experts and novices.

perts and novices, we used a modified t-test developed by Crawford and Howell (1998) for analysis of single-case studies by comparison to small normative samples ($N < 50$). DM's performance on angle judgement tasks did not differ significantly from healthy novice participants. His error was higher than the healthy expert participants for planar angle judgement and for both direct and indirect projected angle judgement, so much so that it trended toward significance ($t(8) = 1.480$, $p = .089$ (one-tailed); $t(8) = 1.526$, $p = .083$ (one-tailed); and, $t(8) = 1.434$, $p = .095$ (one-tailed), respectively). On questions of extent and occlusion, DM's performance was also poor (Figure 5.6). His judgement of extent from an indirect viewpoint was significantly worse than the healthy experts, $t(8) = -2.287$, $p < .05$ (one-tailed), and novices, $t(25) = -1.769$, $p < .05$ (one-tailed).

Particularly because DM was confined to his bed and had no use of his left hand, it was not practical or comfortable for him to perform the direct and indirect drawing tasks. Observing that he still drew often at his bedtable, using small pages (approximately four

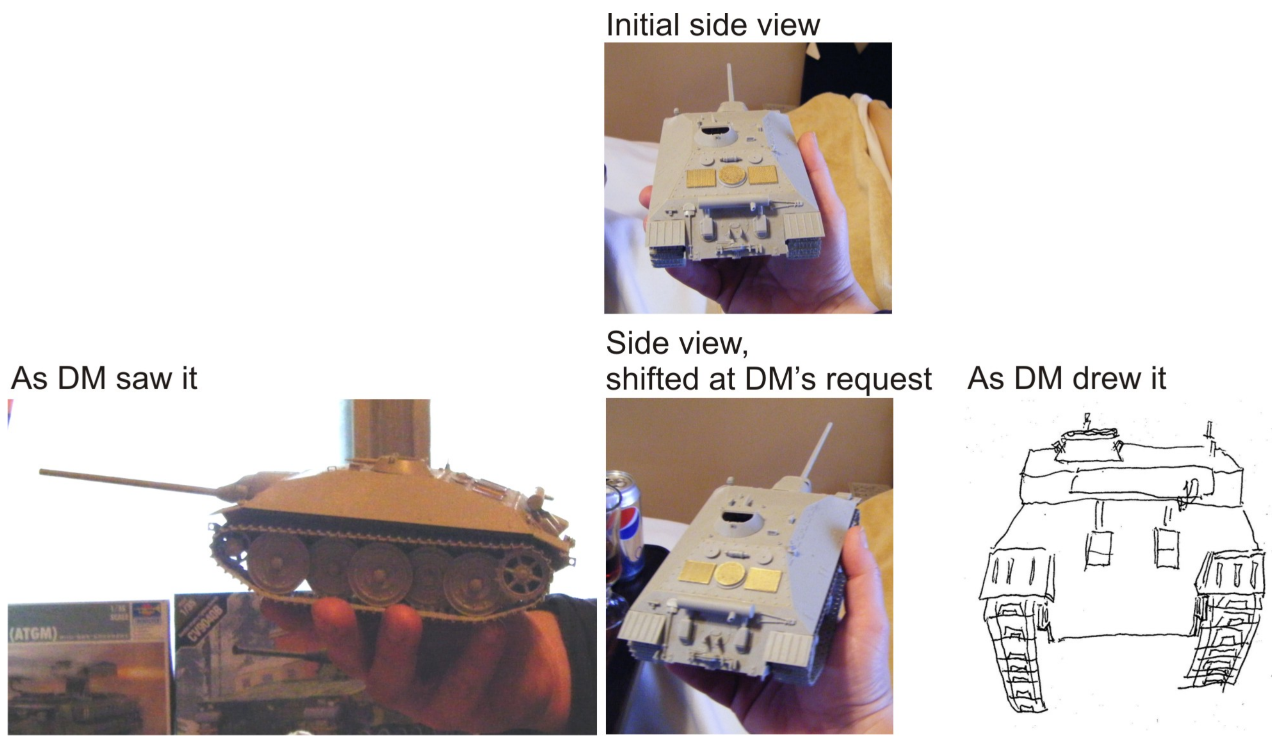

Figure 5.7: Given a side view of a familiar tank (left), DM was asked to draw an imagined side view as seen by the researcher (centre, top). He asked to have the the tank repositioned (centre, bottom), then drew the tank (right).

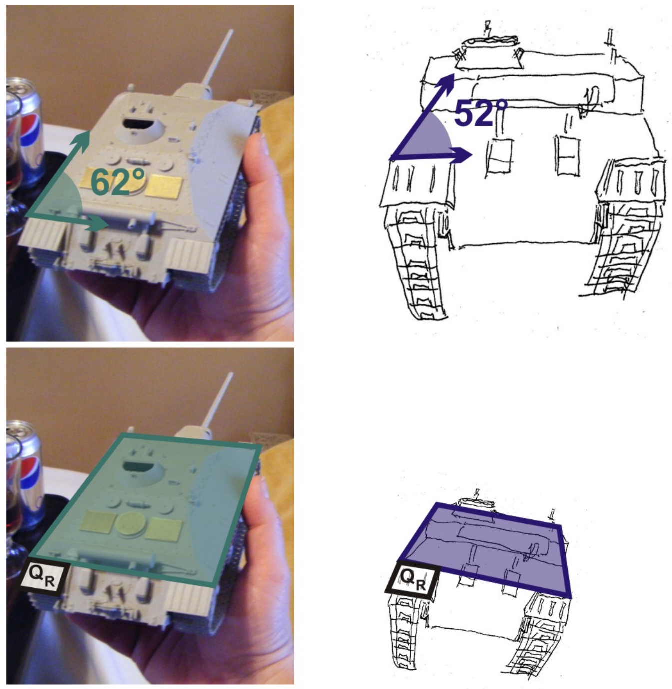

Figure 5.8: Angle drawing accuracy on the corner of the tank (upper). Polygon drawing accuracy of the top surface, scaled to the tread flap polygon Q_R (lower).

inches by six inches) that he could control without the use of his left hand, we improvised an indirect drawing task that accommodated his disability and energy level. We asked DM to draw on a small piece of paper in his usual way, taking as his subject a familiar model tank that he kept in his room, but to draw it as he imagined it would look to the researcher holding it up. (Figure 5.7). He asked for a small adjustment to the angle of the tank so he could glimpse the details of the rear, then drew freely and without time limit.

We measured the lower left corner of the tank (Figure 5.8, upper) in the ground truth photograph (62°) and in DM's drawing (52°). This was the most salient angle in the image but its small error misrepresents the accuracy of the image. The construction of the tank was so far off from the ground truth that there were very few other angles that mapped unambiguously from the drawing onto the ground truth. (For example, the hatch projected—incorrectly—above the tank in the drawing but the gun barrel, a prominent angular feature in the ground truth, barely appeared in the drawing at all.) The percentage magnitude angle drawing error was 16.12%. Expert participants drawing the still life Bravo had a mean percentage magnitude angle drawing error of 26.22% (*SD*= 10.50) on acute angles. If one attributes the misconstruction of the tank to a belief that DM *intended* to draw the tank from straight on (despite asking for the adjusted pose to give him a slightly three-quarters view), then the error is larger. The ground truth angle on the initial side view was 79°, which would produce a percentage magnitude angle drawing error of 34.17%.

The inaccuracies of the drawing would be more clearly described by a polygon-based approach. We illustrated (but did not calculate) the polygon-based accuracy by scaling and rotating the drawing to match the rectangular flap over the tank tread on the left, a suitable reference ROI (Figure 5.8, lower). Treating the "roof" of the tank as a polygon and comparing DM's drawing with the ground truth, there are obviously large errors of size and proportionality.

While DM drew, he commented freely and accurately about what the researcher's viewpoint was. The finished drawing shows some subtle deficiencies on the left side (presumably attributable to neglect) and despite DM demonstrating explicit awareness of the indirect viewpoint, his drawing does not reproduce either the overhead viewpoint or the angular shift he asked for (a failure of mental rotation). The drawing notably distorts the three-dimensional form of the tank, including showing the top of the tank "popping" straight up rather than having a shallow bevel, and by pulling down the treads.

DM's drawings from before his injury show no such distortions of 3D space (Figure 5.4). His performance on the perceptual judgements and (modified) indirect drawing task are consistent with our speculation that drawing requires an accurate mental model of 3D space in multiple reference frames, a model that is clearly disrupted in DM.

Participants like DM—with drawing expertise and a right hemisphere lesion, and yet healthy enough to contribute to the research—are rare. A stronger test of the hypothesis would be to use repetitive transcranial magnetic stimulation (rTMS), a virtual lesion method (Wassermann, 1998), to investigate the role of specific brain areas in drawing. Fierro and colleagues (2000) have previously induced contralateral neglect on a line judgement task by applying rTMS to the right parietal cortex in healthy individuals. Other rTMS studies have shown that the left parietal lobe was important to the generation of mental images but that the right parietal lobe contributed to spatial comparison of those images (Sack, Camprodon, Pascual-Leone, & Goebel, 2005) and that the right parietal lobe was critical to judging object orientation but not identity (Harris, Benito, Ruzzoli, & Miniussi, 2008). These findings are consistent with DM's performance, and we hypothesize that a comparison of healthy age-matched artists and controls drawing from an imagined side view before and during rTMS to the right parietal lobe would reveal measurable reduction in the accuracy of three-dimensionality in the drawing, especially for the artists.

Chapter 6

General discussion

Drawing is a universal human activity and some experts meet its challenges especially well. When we draw, we take in what we see and what we know about the subject, then make marks on the page in a cascade of finely-coordinated motor movements. We inspect those marks while we work, both to correct what has been finished and to plan what is to come. We compare our drawings to the subject and to our intentions. Drawings may be flat, but they create a pictorial space we read as having depth, not just an illusion of volume in objects, but also in the sense of some marks appearing in front of others.

This body of work was initially motivated by a practical professional interest in how we learn to draw. The study of drawing depends upon the study of visual perception and it seems likely that drawing can repay that debt with contributions to vision research. The studies in this thesis developed and validated two objective measures of drawing accuracy, both of which report accuracy locally, as well as for the drawing overall. The measures are consistent with the drawing experience of the people who made the drawings, and with the quality ratings of human critics.

In authentic still life drawing tasks, people drew obtuse angles more accurately than acute angles. The worst-drawn angle represented the corner of a book but a nearby angle of similar size was much more accurately drawn, providing evidence of the top-down influence of context on drawing accuracy. Drawing a book is harder, it seems, than the sum of its angular parts. By comparing angle-based drawing accuracy with performance on several novel perceptual tasks, we showed that skilled drawers have a better model of 3D space. Drawing from an imagined side view demanded mental rotation skill, but overall drawing accuracy did not.

Art teachers are famously pre-occupied with the negative space in drawings. Polygon-based accuracy analysis showed that the big difference between people's drawing in the positive and negative space is in the proportionality and position. We "get the shapes right" and in the right place more in the positive space than in the negative space, even though scaling errors are similar in both and orientation error is higher in the positive space. This is a confirmation for drawing teachers that there's room for improvement in the negative space. It may also signpost the properties of images that matter most to viewers in the figure (rather than ground). We are relatively insensitive to scale between the negative and positive space, but proportions and position in the positive space are special.

Overall polygon accuracy was consistent with ratings of drawing quality, but we were also able to tease apart the dimensions of accuracy in the polygon analysis to discover that mean scaling error was a good predictor of quality. This is an interesting contrast with the positive/negative space finding. When drawing, we do a better job of proportionality and position in the positive space and we are more accurate at orientation in the negative space, but when critics assess overall quality of a drawing, they seem to be influenced by accuracy of *sizes* of shapes. That could mean that we look at "the big picture" differently than we look at the elements within it, that people set different standards when making drawings than when we assess them (although Cohen and Bennett (1997) would suggest otherwise), or, perhaps, that the drawers' extra attention to other properties satisfies the critics and exposes scaling error to judgement.

6.1 Errors of invention

In most of the drawings collected, the viewer could see places where plausible features were locally distorted, such as an edge stretched to fill the space between mis-positioned objects. Such errors were objectively measured by both our angle accuracy and polygon-based analyses. However, inspection of the drawings also revealed a surprisingly common behaviour that was not captured, let alone characterised, by either accuracy analysis. Our analyses could only take into account the *features that had analogues in the ground truth photograph.* We can measure accuracy. We can report absence. We did not anticipate the need to analyse how frequently or dramatically participants drew features that do not appear in the ground truth at all. We describe here an interesting type of error. We saw no mischief or malice, but report a category of frequent natural error that was not yet satisfactorily captured by our accuracy measures: errors of invention.

Our polygon-based analysis is analogous to measuring image deformation, an important research area in computer vision, (Zitov & Flusser, 2003). The study of image deformation includes how to register images to a ground truth—such as overlaying satellite data onto aerial photographs of terrain—and how to reconcile findings from different imaging technologies—such as aligning MRIs with x-rays—despite noise, artifacts, omissions and systematic distortions. The most important difference between the analysis of drawings and general image deformation may be that machines are more constrained by the ground truth than people are, or, as one collaborator said, "People make shit up."

There was clear evidence of errors of invention in the drawings of still life Bravo. If a participant studied the still life from the correct viewpoint, the feet on the left side of the stool were both occluded by the tin of tools. Despite this, 25 of the 35 participants (71.4%) drew one left foot, or both, peeking out from behind the tin of tools (e.g., Figure 6.1). This particular error of invention was captured, in a small way,[1] by the overall angle accuracy of the drawing, but we discuss it here as an unambiguous and commonplace example of invention in the dataset.

The "peekaboo foot" could simply be symptomatic of a lapse of viewpoint; a participant might glimpse the foot when leaning too far left. However, we note that the study was conducted in a small room where both the still life table and the participants were positioned very close to a wall on the left. In fact, the first left-handed participant complained that he had no elbow room. We arranged a little more clearance for all subsequent participants, but not much. It would have been very difficult for the participants to lean far enough left to see the "peekaboo foot," and impossible for them to do so in comfort.

Boundary extension has previously been observed in memory drawing tasks (Intraub & Richardson, 1989) where participants included more of the background than was visible in the stimulus. The "peekaboo foot" error differed from boundary extension in several ways. It was not on the periphery of the image (and participants were not limited to a specific image size or format), it was not a simple filling-in of more background detail or texture, and it was not easily attributable to memory since participants had a clear view of the subject the whole time they were drawing.

We have no metric for the size of the mistake, just its presence. Analysis showed a relationship between making the "peekaboo foot" mistake and a participant's years of drawing experience that approached significance ($F(1, 33) = 3.619$, $p = .066$, $\eta^2 = .099$) but no relationship to drawing accuracy at all ($F(1, 33) = .025$, $p = .874$, $\eta^2 = .001$). Participants who made the "peekaboo foot" mistake had fewer years of drawing experience ($M = 6.80$, $SD = 8.31$) than the participants who avoided it ($M = 15.80$, $SD = 20.04$) but

[1]When inventing the "peekaboo foot," participants usually altered the relationship between the edge of the gray mat and the side of the tin of tools, an alteration that would be captured in at least one angle measure.

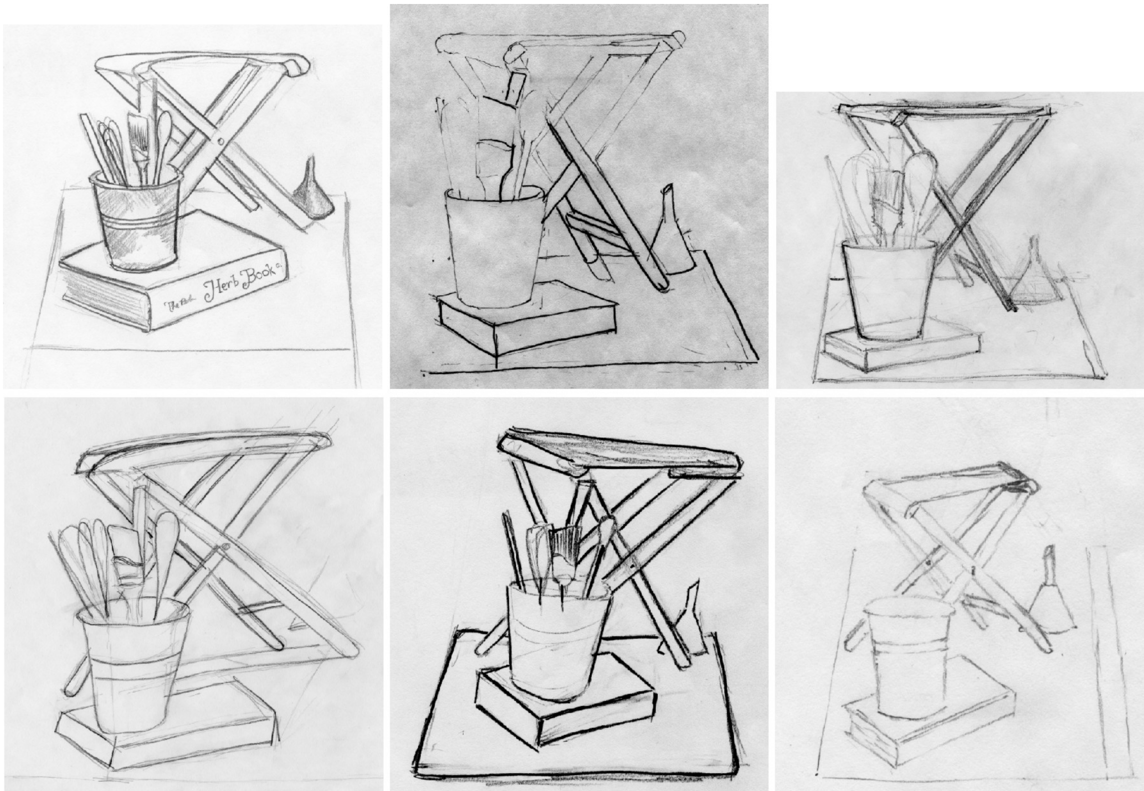

Figure 6.1: In the upper row, drawings correctly show the tin of tools occluding the left feet of the stool. In the lower row, participants have all made the "peekaboo foot" mistake, stretching or shifting the image of the stool such that its left foot peeks out beyond the tin of tools.

the angle-drawing accuracy had little chance of detecting it (Figure 6.2).

The "peekaboo foot" mistake represents a distinct new topic of drawing analysis. To study it further, we need to extend the accuracy measures to detect and measure errors of invention.

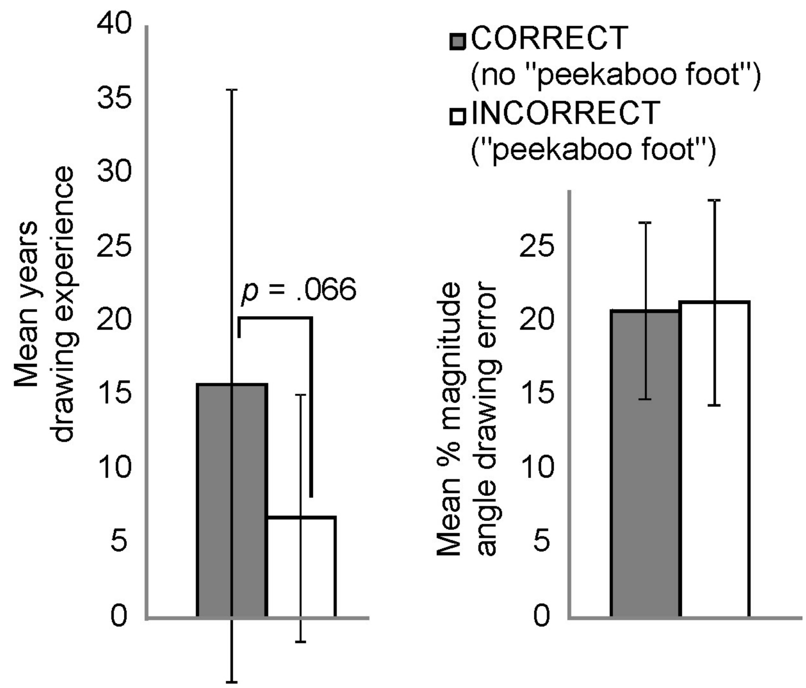

Figure 6.2: Years of drawing experience as a function of whether the drawer made the "peekaboo foot" mistake (left) and drawing error as a function of whether the drawer made the "peekaboo foot" mistake (right). Error bars represent ± one standard deviation.

6.2 A proposal to study corrections and drawings in progress

Once we can measure error in a drawing, the inevitable question is, *When and where does error arise?* The popularization of tablet computers opens up a new opportunity for polygon-based accuracy measurement and the direct capture of the time course of a drawing. It is now practical to assess the interim states of drawings, which have not been well-studied. One of the most interesting and frustrating things about drawing is how early mistakes have lingering consequences (van Sommers, 1984). Research has only rarely examined drawing in a dynamic fashion (Kozbelt & Serafin, 2009). We know that artists regularly plan and evaluate work in progress, and engage in metacognition about emerging drawing (Fayena-Tawil, Kozbelt, & Sitaras, 2011). We propose, in future work, to study the development of representational drawings over time. Rating completed drawings does not approximate the drawer's ecological experience of the work in progress. Sometimes a rating of drawing quality is a proxy for the real property of interest, whether drawers detect and attempt to correct errors in their drawings, which we can detect and measure with a tablet and polygon-based analysis.

We are particularly interested in corrections: when and where people detect problems in their drawings, and how they attempt to improve the drawing. This is an important intersection of drawing and visual perception. What do we notice, and when do we notice it? In an influential study, *Why Can't Most People Draw What They See?,* Cohen and Bennett (1997) broke the drawing process down into four parts: perceiving the subject; deciding what to represent and how; making the physical marks on the page; and, assessing the drawing in order to make any necessary corrections. They looked for evidence of the extent to which drawing error could be attributed to specific parts of the process in a series of four experiments, *all reliant on human critics for ratings of drawing accuracy.* Based on their first three experiments, the authors argued persuasively that, since people had enough skill to trace images and made good decisions about where to make marks when tracing, the likeliest or greatest sources of drawing error are presumably in the visual perception of the subject and in assessing a drawing in progress to make corrections. As a probe of the latter, they tested the hypothesis that drawers misjudge the accuracy of their own (completed) drawings. Comparing drawing ratings made by the drawers and other people, Cohen and Bennett found no evidence that we are particularly blinded to our own drawing errors after the fact. However, that is not the only way that people's perception of their own drawings would influence accuracy. We don't make drawings all at once or all over the page. We make drawings one mark at a time, we make mistakes one mark at a time, and we make corrections one mark at a time. Furthermore, when we look at our drawing in progress, we may rejoice or despair about the accuracy globally, but we need to detect and repair errors locally. The assessment stage of the drawing process is multi-faceted. In order to make corrections, artists need to detect the presence *and location* of errors. It seems worthwhile to revisit the question of whether and how our perception of the drawing contributes to drawing accuracy by applying objective accuracy measures and studying drawings in progress.

In a famous riddle, two brothers once competed in a peculiar camel race where the winner would be the one who finished second. At the rate they set out, the race might have taken an eternity, but someone gave them this advice: Swap camels. To measure whether people's experience of working on their own drawing biases their assessment of it, we would propose to have them inspect *unfinished* drawings (theirs, and others), highlight errors and make corrections to the best of their ability.

6.3 Other future work

One of the prosaic contributions of these studies was the development of a practical, ecological still life drawing task. First, the still lifes were layered and complex to ensure that the experience was authentic and to avoid a ceiling effect with expert drawers. We marked objects unobtrusively so they could be consistently placed, we affixed objects to one another to ensure that slant and projected angles were maintained, we set up landmarks for aligning the participant with the same viewpoint every time and we measured the ground truth with photographs, clinometer and protractor. Now that we have some sound observations as the basis of future work, still lifes should be purpose-built to test specific hypotheses. For example, to investigate the bird's eye error, while we could apply a polygon-based analysis to the indirect drawings already collected, we would prefer to build a new still life, incorporating surfaces a priori at, above and below eye level. *How far above eye level does a surface have to be before we stop showing it from overhead?*

Previc and Intraub (1997) argued that vertical biases in many visuomotor behaviours were a consequence of the likelihood that the lower visual field represents near space and the upper visual field usually represents far space. The bird's eye view error is an example of such a vertical bias. The authors suggested that neural systems mediating operations in peripersonal space (such as reaching) would be biased toward the lower visual field (which usually represents near space) and those mediating operations in extrapersonal space (such as object recognition) would be biased toward the upper visual field (which usually represents far space). In drawings of photographs from memory, they found a tendency to shrink the objects in the image and shift them lower in the scene, but the effect was reversed for a few photographs such as one showing a fruit plate from overhead (already a bird's eye view). The authors proposed two hypotheses to account for their findings: that participants recall the scene as more distant, or recall it as seen from a more overhead view. Based on their evidence, it is difficult to tease the two apart. Their photograph-based task unavoidably confounds pictorial depth (near/far) with height in the visual field. That confound could be eliminated by drawing from a still life or other three dimensional subject with objects of interest at predetermined distances and controlled heights in the visual field. *Are distortions of drawing from memory more characteristic of a shift in viewpoint from near to far or from low to high?*

Shape constancy, the useful ability to recognize familiar objects even when presented from unfamiliar viewpoints, can have an unfortunate effect on drawing. When presented with an oblique view of a circle, people draw neither a circle nor the elliptical retinal projection, but something in between (Thouless, 1931). Furthermore, such drawing errors are affected by top-down knowledge about the subject (Taylor & Mitchell, 1997) and perhaps even by what Picard and Durand (2005) and Matthews and Adams (2008) describe as *"idiosyncratic canonical bias"*, a tendency to distort observational drawings toward one's habitual distortions in free drawing of a similar subject. Cohen and Jones (2008) suggested that drawing skill is accompanied by an ability to overcome shape constancy, though McManus and colleagues (2011) failed to replicate those findings and Perdreau and Cavanagh (2011) found no evidence that experience gives artists a perceptual advantage. This work has depended on simple stimuli like circles and cylinders, presumably chosen, at least in part, to be easy to measure. However, the experience of shape constancy may not affect more complex stimuli uniformly. Polygon-based accuracy measures could be used to discriminate between objects whose silhouette was drawn accurately and those whose interior features were well-depicted and well-placed. *To what extent are drawing errors attributable to flaws in the silhouette of an object, to flaws in interior features of the object, or to both?*

The accuracy measures in this thesis were originally designed to study the making of drawings, but they can obviously contribute to the study of the perception of drawings and other images. Object recognition is increasingly understood to be a product of multiple systems which rely on different types of stored information, visual and semantic, viewpoint-specific and viewpoint-independent (Logothetis & Sheinberg, 1996). *Image-based theories* suggest that object recognition relies on transforming the overall view of the object to match an internal view (Ullman, 1998). *Structural description theories* suggest that object recognition relies on observation of features and relationships between them (Biederman, 1987). Polygon-based accuracy measures could be used to model how well people recognize objects with controlled amounts of distortion to the silhouette—which could isolate image-based recognition—and interior features—which could isolate structural descriptors. *How robust is object recognition to distortions of the silhouette and interior features?*

Since ratings of general quality correlate with accuracy, future work can contribute to unpacking the *dimensions* of quality. First, in a similar study, we can explicitly focus on people's ratings of drawing accuracy, and compare those results to our objective accuracy measures. *How accurate are human ratings of drawing accuracy?*

We can compare objectively measured accuracy to people's ratings of other dimensions of quality, such as beauty and originality (as in Kozbelt, 2006). *Which other aspects of drawing quality correlate with accuracy?*

While accuracy and quality were generally consistent with one another, there were some drawings that were rated of surprisingly high quality given their accuracy. These drawings may have properties of mark-making—what is often called "line quality" in the studio—that make them look better, or worse, than we'd expect from their accuracy. Some of the most highly-rated drawings may have stood out because they included areas of tone (such as the dark hatching in the background of the highest-rated drawing in Alpha, Figure 4.7, upper left). Using the polygon-based accuracy measures of Chapter 3, we can artificially alter copies of these surprising drawings to incorporate controlled amounts of error, to add or remove tonal areas, or to mirror the mark placement of a differently-rated drawing. Thus, we can begin to characterize "goodness" in mark-making. *Which perspective drawings of the same accuracy but different marks are judged to be higher quality?*

Anecdotally, it seemed to be the case that novice artists recorded details (such as the title on the spine of the book or the measurements on the ruler) more than experienced artists. Furthermore, the most highly-rated drawing of Bravo (Appendix B.2) included such details but was of only middling accuracy. We can artificially alter these drawings to remove some details. *Are drawings of equal accuracy rated more highly when they include details?*

As we proposed in Chapter 3, we can construct drawings of controlled amounts of error of different types and invite human critics to rate their accuracy. We can contribute to a better understanding of visual perception by asking, *How sensitive are people to errors of orientation, proportionality, scale and position in a drawing?*

As we also proposed, in Chapter 5, we could use rTMS to study the influence of the dorsal stream on the mental model of three-dimensional space in artists. With these drawing accuracy measures we could clarify whether the main contribution is made by the dorsal stream itself, or by communication between the dorsal and ventral streams. *Does rTMS of the right superior or inferior parietal lobule impair drawing from an imagined side view in artists and controls, and are artists affected more?*

Snyder (2009) proposes that savant skills in the realm of drawing are effectively the result of failures of top-down processing: a savant draws in extraordinary detail and accuracy because she or he is not influenced as heavily as most people by the context of the overall image. Snyder and colleagues (2003) investigated this hypothesis by interfering with left fronto-temporal lobe function in healthy individuals with rTMS to see if it would trigger savant-like changes in drawing and proofreading tasks. In drawing from memory and from imagination, there were some stylistic changes but no systematic improvement. It could be argued, though, that the drawing tasks were poorly suited for isolating the influence of top-down processing. If the explanation of savant drawing skill is sound and the left fronto-temporal lobe is an appropriate neural correlate of top-down cognitive processing of the scene, then on a still life drawing task such as in Chapter 2, under the influence of rTMS, the corner of the book should be drawn as accurately as any other similarly-sized angle in the image. *Are healthy individuals with virtual left fronto-temporal lesions from* *rTMS equally accurate when drawing small angles isolated in an image and found on the receding surface of a boxy object?*

6.4 On measuring drawing accuracy

Researching drawing based on the ratings of human critics is an informative and reasonable approach. Drawing training and appreciation are also based on the critic's eye. But the critic's eye is subjective and variable, and sometimes it is hard to be certain just what is being rated. Our main motivations for developing the use of objective measures of accuracy are the possibility that human raters compound or confound the very perceptual biases of drawers that we hope to detect and study, that human raters use too coarse a scale to be able to detect differences that are present, and that human raters may be conflating several dimensions of drawing error into one rating, dimensions we might be able to unpack with measured accuracy (e.g., local versus global accuracy, as in Chapter 2, and errors of orientation, proportionality, scaling and position, as in Chapter 3). Our accuracy measures promise to be more precise, more objective, more consistent and more clearly defined than most ratings commonly used in the extant literature. It is clear what we're measuring and we can discriminate between different measurable properties. Though this thesis focused on perspective drawing from the observation of three-dimensional still lifes, both of the accuracy measures could be applied equally well to drawings from photographs, tracing or copying, i.e., any drawing task for which there is a known ground truth.

In the angle-based accuracy study in Chapter 2, a tally of missing angles was meant to serve as a simple measure of drawing completion. An unfinished drawing would obviously have fewer complete angles to analyse, but some bad drawings would also have too few angles. When objects are drawn seriously out of position, the angles that should have been created by their overlap with other features would be lost. The highest- and lowest-rated drawings in Bravo for quality had nearly identical mean percentage magnitude angle drawing error (Figure 4.7). The bad drawings were few enough—or good enough—that they did not obscure a significant relationship between quality and accuracy, but outliers should be excluded in future analyses. We would propose to set an appropriate threshold, some minimum proportion of angles that should be present in the drawing for mean percentage magnitude angle drawing error to be meaningful. The polygon-based approach minimized this problem by distinguishing missing landmarks and outright errors of occlusion.

The angle-based approach (of Chapter 2) is low-tech and fast. It is most appropriate for complete, sound drawings, for drawings of subjects made up mainly of straight lines, for the analysis of local error within the drawing, and especially for angle-specific research questions such as the study of divergent perspective. Simple, incomplete or deeply flawed drawings would be poorly-described by the overall angle accuracy measure because too few angles would contribute to the mean.

The polygon-based approach (of Chapter 3) takes more time and computation but it has several advantages. It appears to be more robust to, and informative about, incomplete drawings. Only complete polygons are analysed and contribute to the overall measure of polygon accuracy. The polygon-based approach is suitable for extending accuracy measurement to curvilinear subjects, which would broaden the reach of drawing research to the full range of subject matter, e.g., landscape, portraits, and the figure. The model can include angle accuracy when desirable, and is designed to adapt to any research question for which there are informative, unambiguous landmarks in the ground truth. Because it uncouples errors of object position from errors of proportionality, orientation and scaling, the polygon-based measure of drawing accuracy is particularly well-suited to investigate the relationship between drawing the visual properties of objects—presumably related to ventral stream functions of object identification—and locating the object correctly in the pictorial space—which may have more to do with dorsal stream functions about where objects are located (Mishkin et al., 1983) or how we act on them (Milner & Goodale, 2006) or how we navigate toward them (Schindler et al., 2004).

What does this mean for training artists? It is conventional wisdom in the teaching studio that artists must practise drawing by direct observation of 3D subjects, not just from copying photographs and other drawings. In these studies, some of the few or first to actually ask people to draw from 3D subjects, the findings show that the people who draw most accurately are better judges of slant and spatial extent, suggesting that they have a better mental model of the 3D properties of the still life. The unanswered question is causality. This work alone cannot determine whether drawing practice improves mental models of 3D space, or if the mental model supports more accurate drawing, but it has laid a solid foundation for controlled objective and ecological research into the teaching, learning and practice of drawing.

APPENDICES

Appendix A

Illustrating polygon error calculations

A.1 Orientation

When we compare the orientation of the polygon Q_1 in drawing $n7$ to the orientation of its reference polygon, Q_R, we can see that it is rotated clockwise more than Q_1 is in the ground truth photograph. Informally, we measure the orientation error of polygon Q_1 in a drawing such as $n7$ by:

1. Reorienting the drawing to match the angle of the minor principal axis of its reference polygon, Q_R, to that of the ground truth photograph (middle figures);
2. Taking the angle—in the ground truth photograph—of the minor principal axis of Q_1 to the minor principal axis of Q_R (lower left figure);
3. Taking the angle—in the drawing—of the minor principal axis of Q_1 to the minor principal axis of Q_R (lower right figure); and then
4. Taking the magnitude of the difference between those two angles.

In the ground truth photograph, the angle of Q_1 ($\alpha_{photo,1}$) is -1.4816 radians. In drawing $n7$, the angle of Q_1 ($\alpha_{n7,1}$) is -1.1214 radians. The orientation error of Q_1 in drawing $n7$ is $|\alpha_{photo,1} - \alpha_{n7,1}| = 0.36886$ radians (21 degrees).

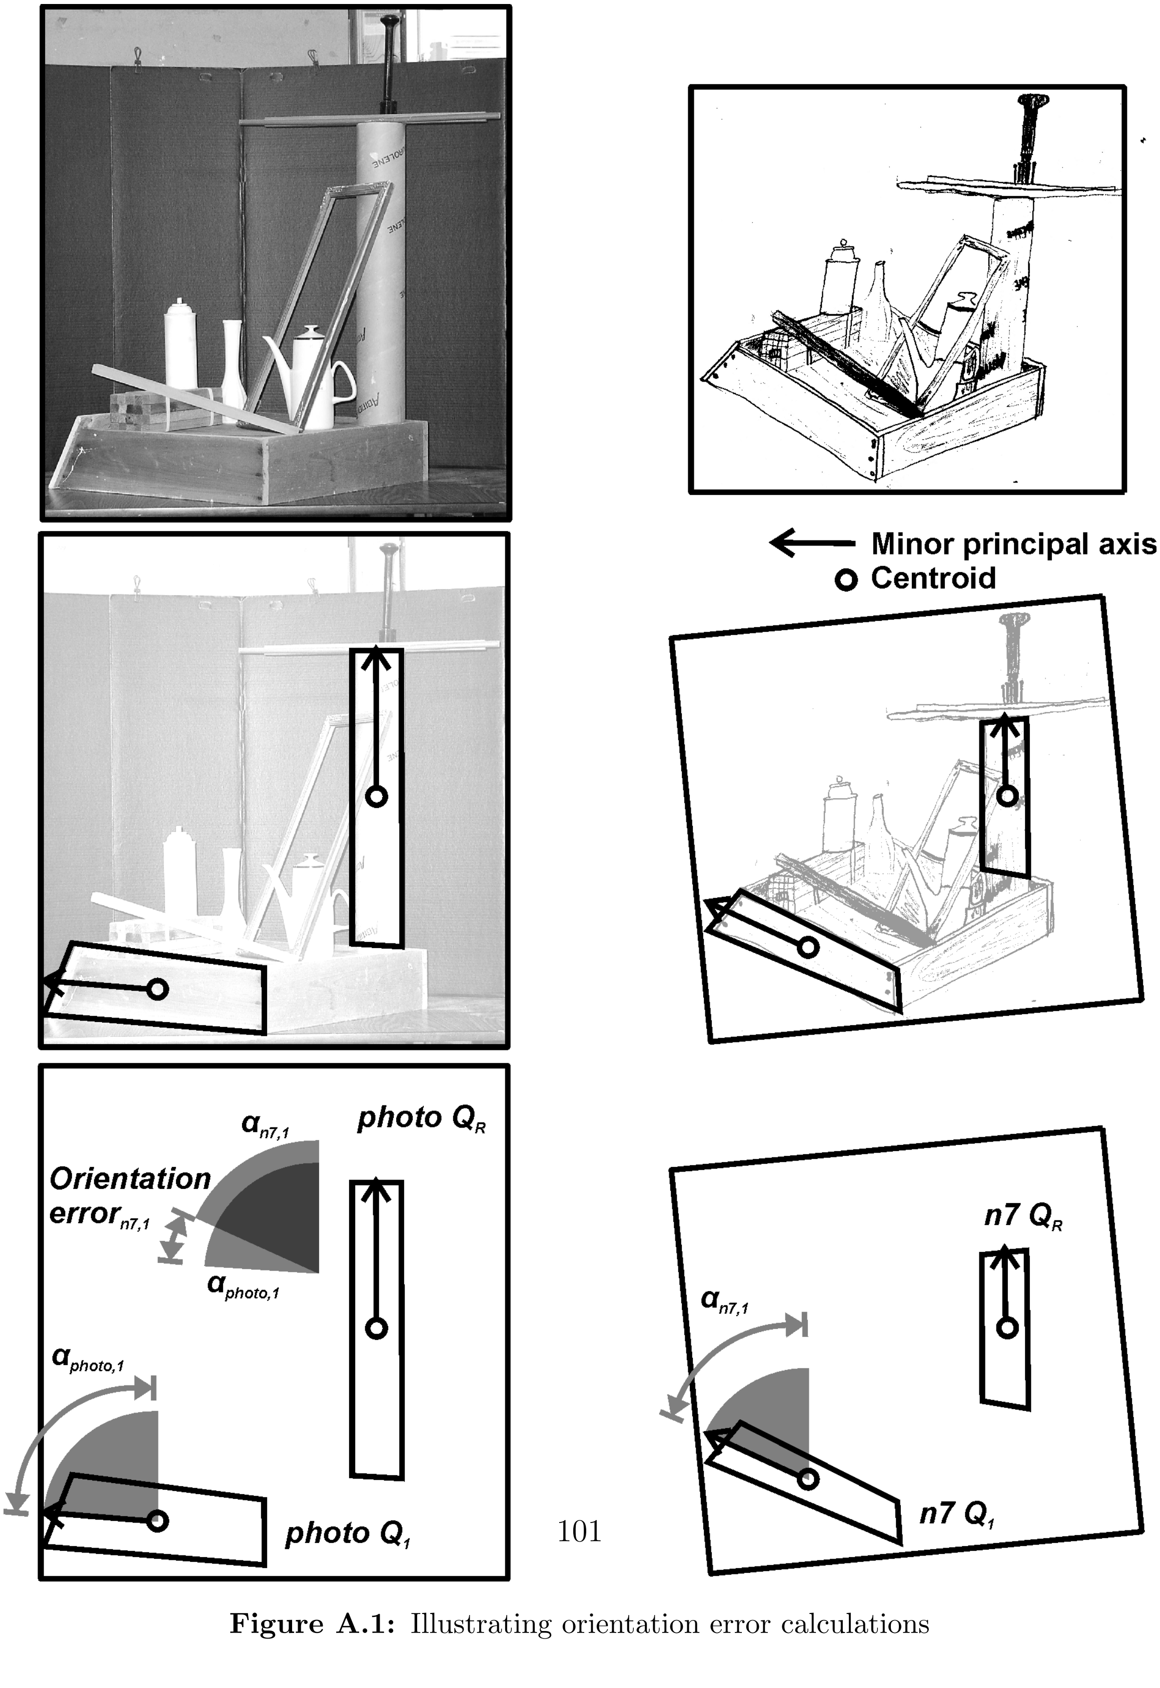

Figure A.1: Illustrating orientation error calculations

A.2 Proportionality

In the ground truth photograph, polygon Q_1 is nearly rectangular and roughly three times as wide as it is tall. In drawing $n7$, polygon Q_1 is longer and skinnier, roughly four times as wide as it is tall, and less rectangular in shape. Informally, we measure the proportionality error of polygon Q_1 by measuring side lengths—to capture shrinkage or stretch—and ray lengths from the centroid—to capture shear. For example, in the ground truth photograph, the left side of Q_1, $L_{photo,4}$, is 14% of the total perimeter (left figure) but in drawing $n7$, $L_{n7,4}$, is only 11% of the total perimeter (right figure), so the side $L_{n7,4}$ contributes an error of 3% to the total perimeter error. Similarly, the ray $R_{photo,3}$ is 26% of the total of the ray lengths (left figure), but in drawing $n7$, ray $R_{n7,3}$ is 28% of the total (right figure), so the ray $R_{n7,3}$ contributes an error of 2% to the total ray length error. We calculate the proportionality error of polygon Q_1 by halving the sum of all these side length and ray length errors. The proportionality error of Q_1 in drawing $n7$ is 0.085369, or 8.5%.

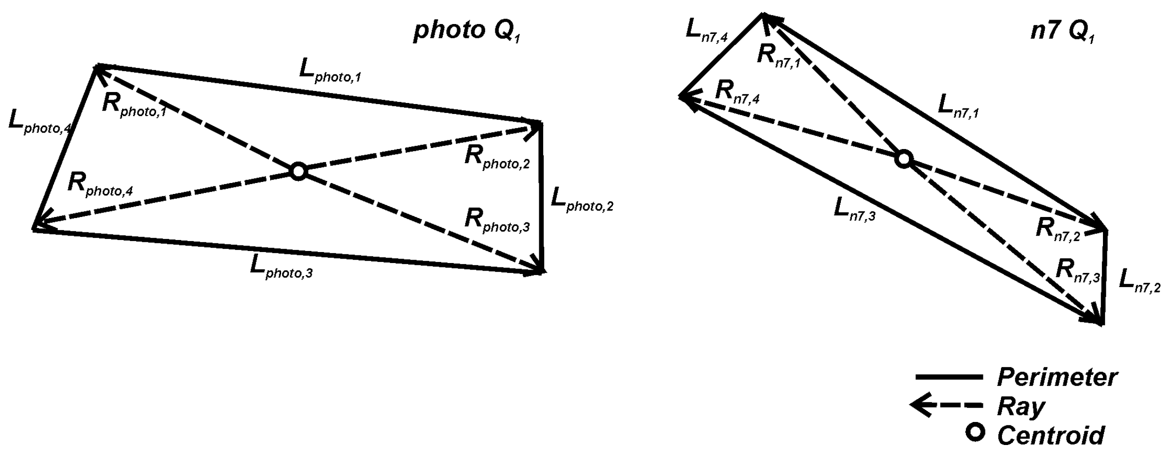

Figure A.2: Illustrating proportionality error calculations

A.3 Scaling

In the ground truth photograph, the area of Q_1 is virtually identical to the area of its reference polygon, Q_R. In drawing $n7$, the area of Q_1 is too large, roughly 124% of the area of its reference polygon, Q_R, while in drawing $j4$, the area of Q_1 is too small, roughly 41% of the area of its reference polygon Q_R. Informally, we measure the scaling error of polygon Q_1 in a drawing by:

1. Taking the ratio—in that drawing—of the area of Q_1 to the area of its reference polygon, Q_R;
2. Taking the ratio—in the ground truth photograph—of the area of Q_1 to the area of its reference polygon, Q_R;
3. Taking the magnitude of the difference between those ratios; and then
4. To keep the scaling error proportional to the true size of Q_1, dividing by the ratio—in the ground truth photograph—of the area of Q_1 to Q_R.

The scaling error of Q_1 in drawing $n7$ is $\frac{|124\%-100\%|}{100\%} = 24\%$. The scaling error of Q_1 in drawing $j4$ is $\frac{|41\%-100\%|}{100\%} = 59\%$.

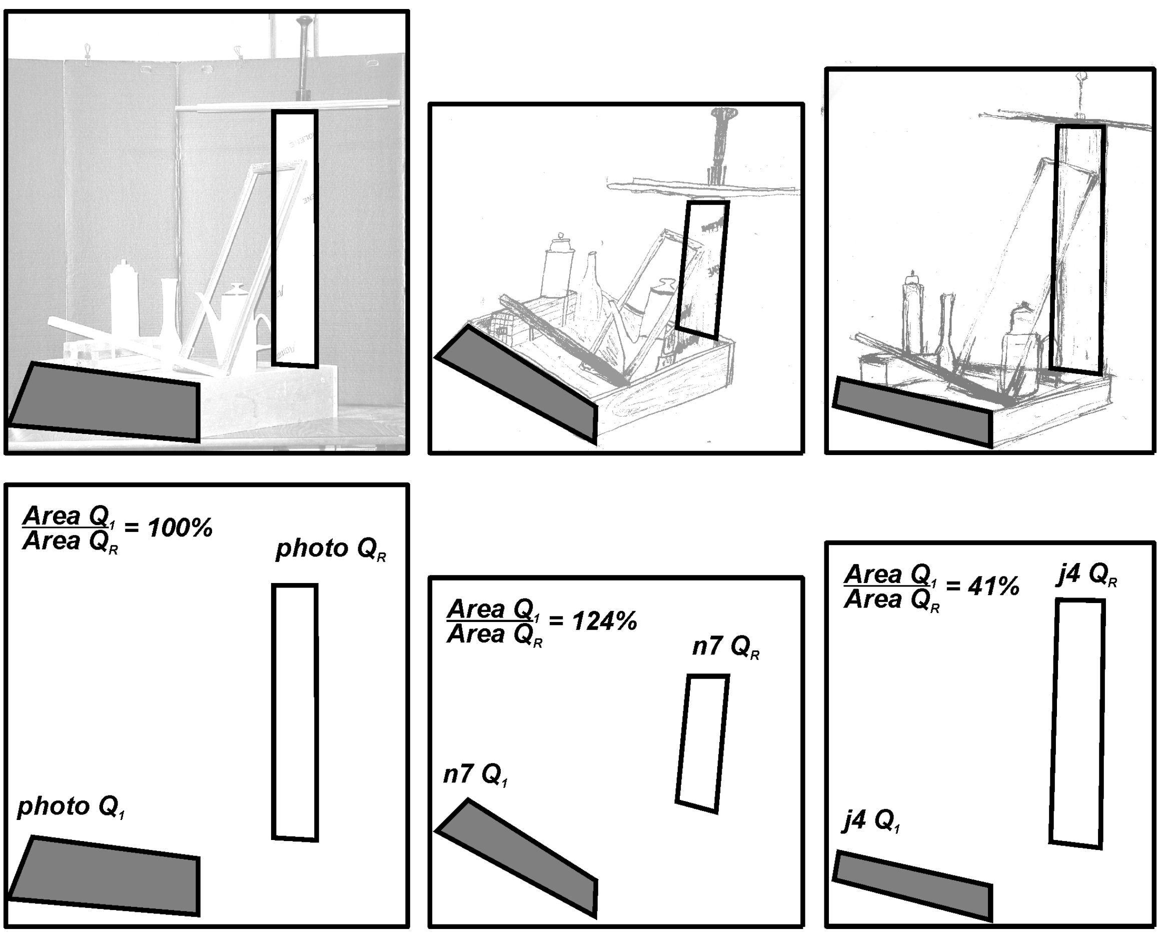

Figure A.3: Illustrating scaling error calculations

A.4 Position

When we resize, reorient and align drawing $n7$ to match its reference polygon, Q_R, to that of the ground truth photograph, we can see that—in drawing $n7$—polygon Q_1 is lower and to the left of where the ground truth photograph shows it should be. Informally, we measure the positioning error of polygon Q_1 in a drawing such as $n7$ by:

1. Resizing the drawing to match the area (A_R) of its reference polygon, Q_R, to that of the ground truth photograph (upper figures);
2. Reorienting the drawing to match the angle of the minor principal axis of its reference polygon, Q_R, to that of the ground truth photograph (middle figures);
3. Aligning the drawing to the ground truth photograph by aligning the centroids of their respective reference polygons, Q_R (lower figures);
4. Taking the distance, D_{offset} , between the centroids of polygon Q_1 in the drawing and in the ground truth photograph (lower figures); and then
5. To keep the positioning error proportional to the true distance from Q_1 to Q_R, dividing D_{offset} by the distance D_{Q_1}—in the ground truth photograph—from the centroid of Q_1 to the centroid of Q_R (lower figures).

Position is the distance from the centroid of the reference polygon, and it is not given in absolute units, e.g., inches or pixels. It is measured in multiples (or fractions) of the square root of the area of the reference polygon, $\sqrt{A_R}$, which would be the height of a square of equal area. This reference length is intentionally dependent on the *area* of the reference polygon—rather than its height—to make position measurement less sensitive to deformations of the shape of the reference polygon

In the ground truth photograph, the centroid of polygon Q_1 is $2.4286\sqrt{A_R}$ from the centroid of its reference polygon, Q_R. In drawing $n7$—resized—the centroid of Q_1 is $3.0685\sqrt{A_R}$ from the centroid of its reference polygon, Q_R. The offset D_{offset} between the centroids of Q_1 in the ground truth and in drawing $n7$ is $0.68011\sqrt{A_R}$. The positioning error of Q_1 in drawing $n7$ is $fracD_{offset}D_{Q_1} = 0.28005$, or 28%.

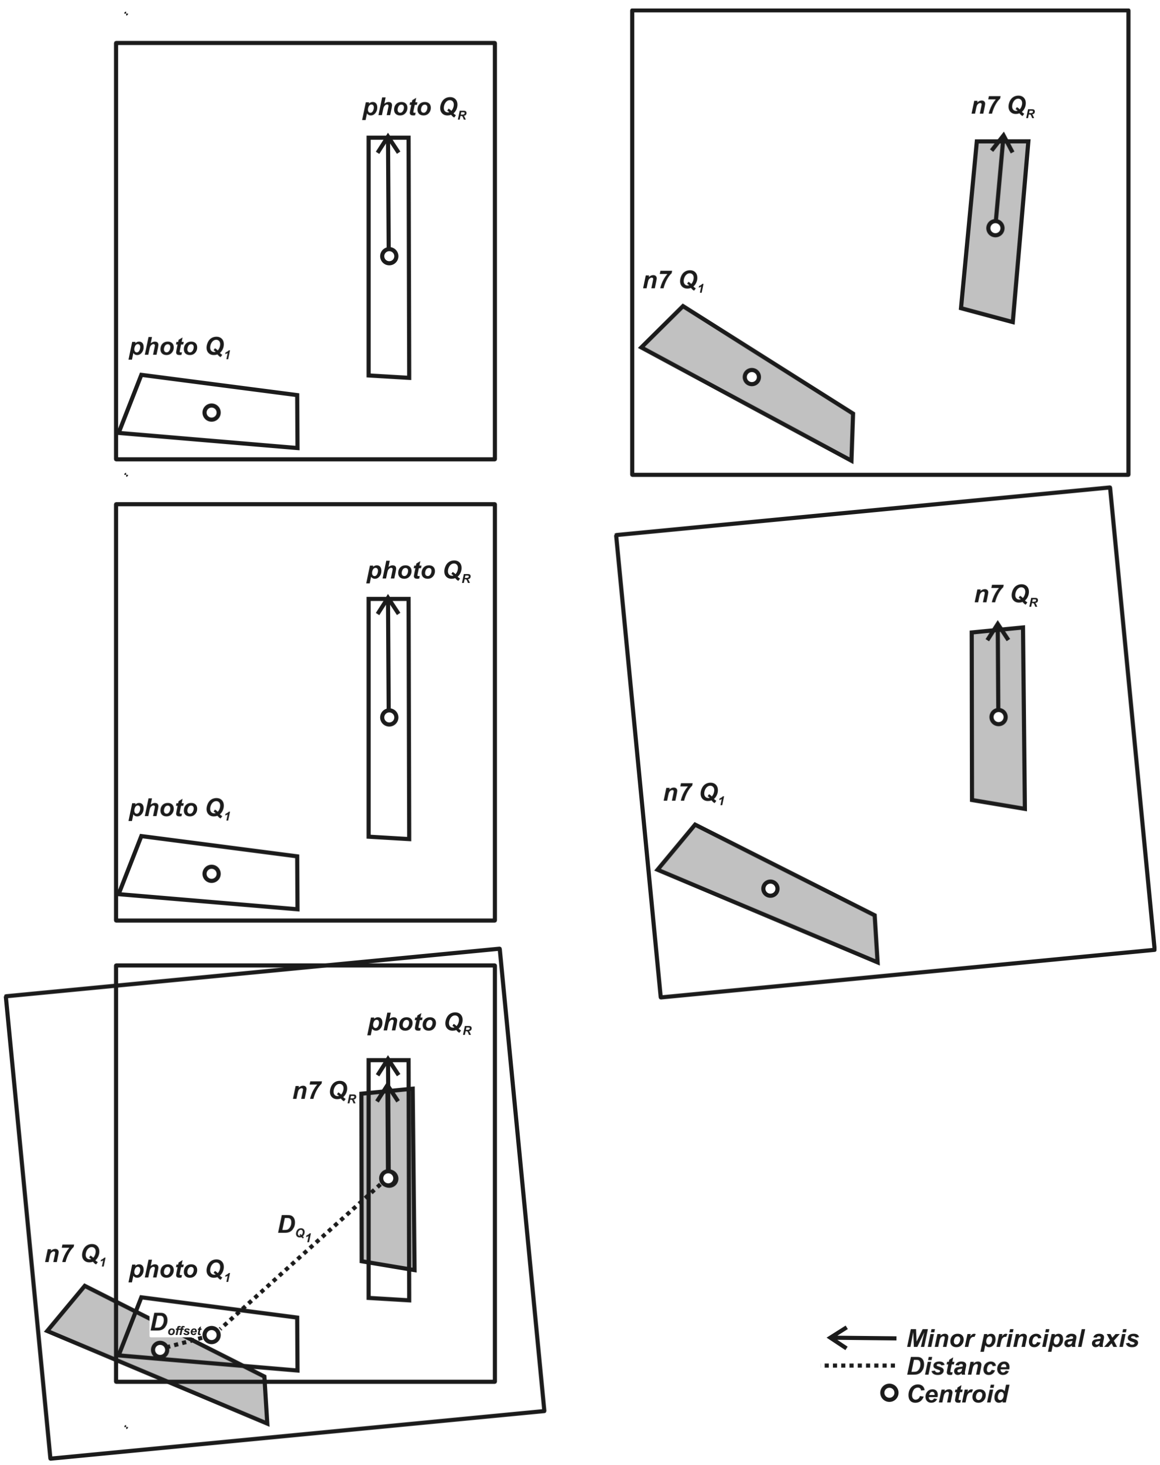

Figure A.4: Illustrating position error calculations

Appendix B

Examples of participant drawings, superimposed on ground truth

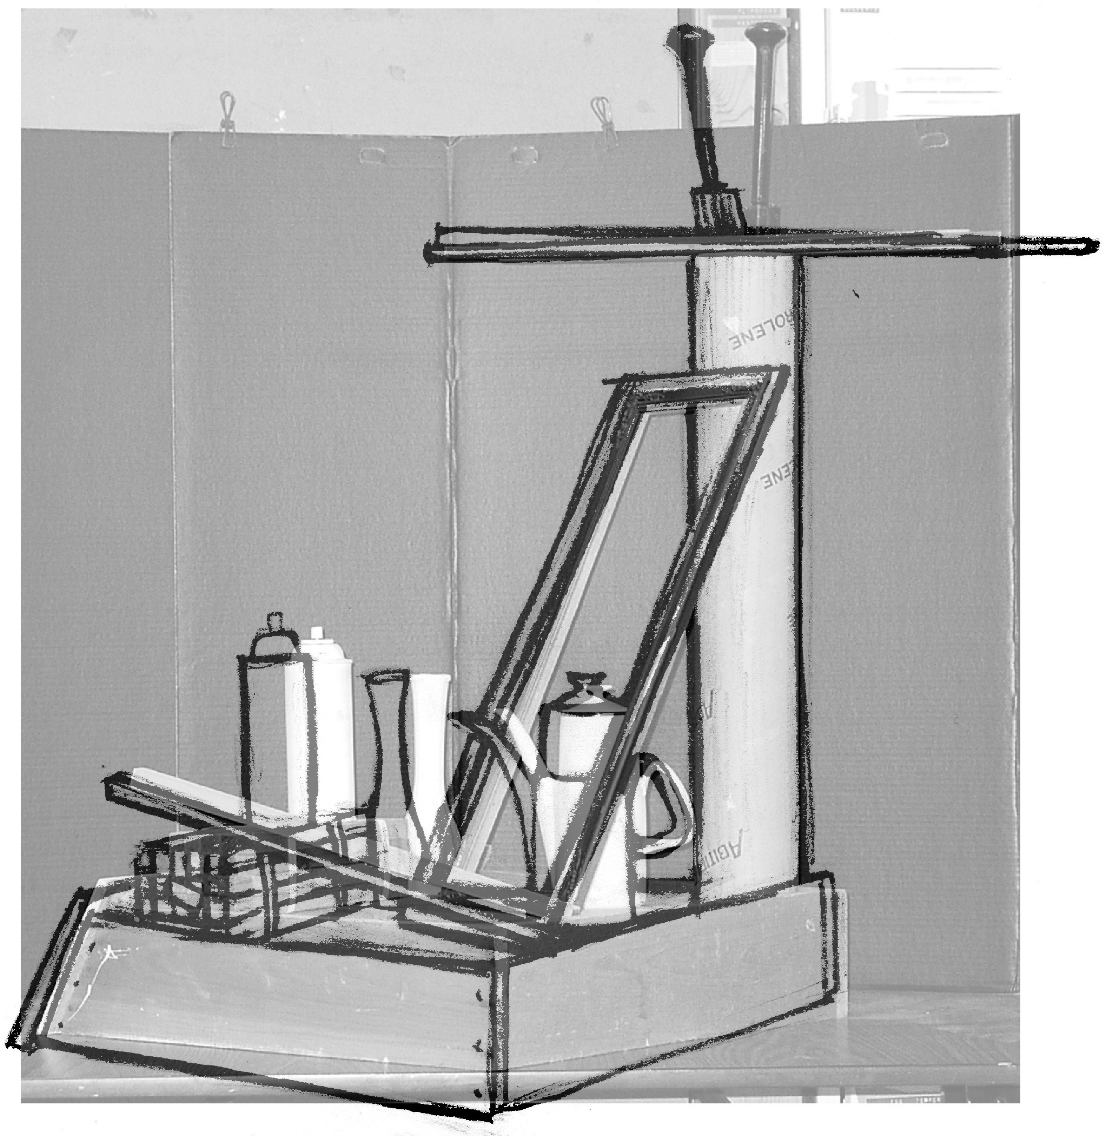

Figure B.1: The drawing with the lowest overall polygon error, $\mathbb{G} = .13$, manually resized and superimposed on the ground truth photograph of still life Alpha, aligning the images of the cylinder on the right (which was also used as the reference polygon). The drawing was made by a novice participant with three years of reported experience. It was rated 5.50 by human raters (which put it 12^{th} in the set of 34 drawings of Alpha).

Figure B.2: The drawing with the lowest mean percentage magnitude angle drawing error, 7.49%, manually resized and superimposed on the ground truth photograph of still life Bravo, aligning the images of the tin of tools on the left. The drawing was made by a novice participant with nine years of reported experience. It was rated 5.28 by human raters (which put it 5^{th} in the set of 35 drawings of Bravo).

Figure B.3: The drawing with the lowest mean percentage magnitude angle drawing error, 8.31%, manually resized and superimposed on the ground truth photograph of still life Charlie (in side view), aligning the images of the mug on the left. The drawing was made by a novice participant with one year of reported experience and a mental rotation task score of 8/24. It was rated 3.38 by human raters (which tied it for 17^{th} in the set of 35 drawings of Charlie).

Appendix C

Age, sex and STEM training

Without making any claims supporting or denying a relationship between mental rotation ability and either drawing accuracy or expertise in our findings of Chapter 2, we report the makeup of the participant pool and its performance on mental rotation by sex, STEM education and age (Figures C.1 and C.2).

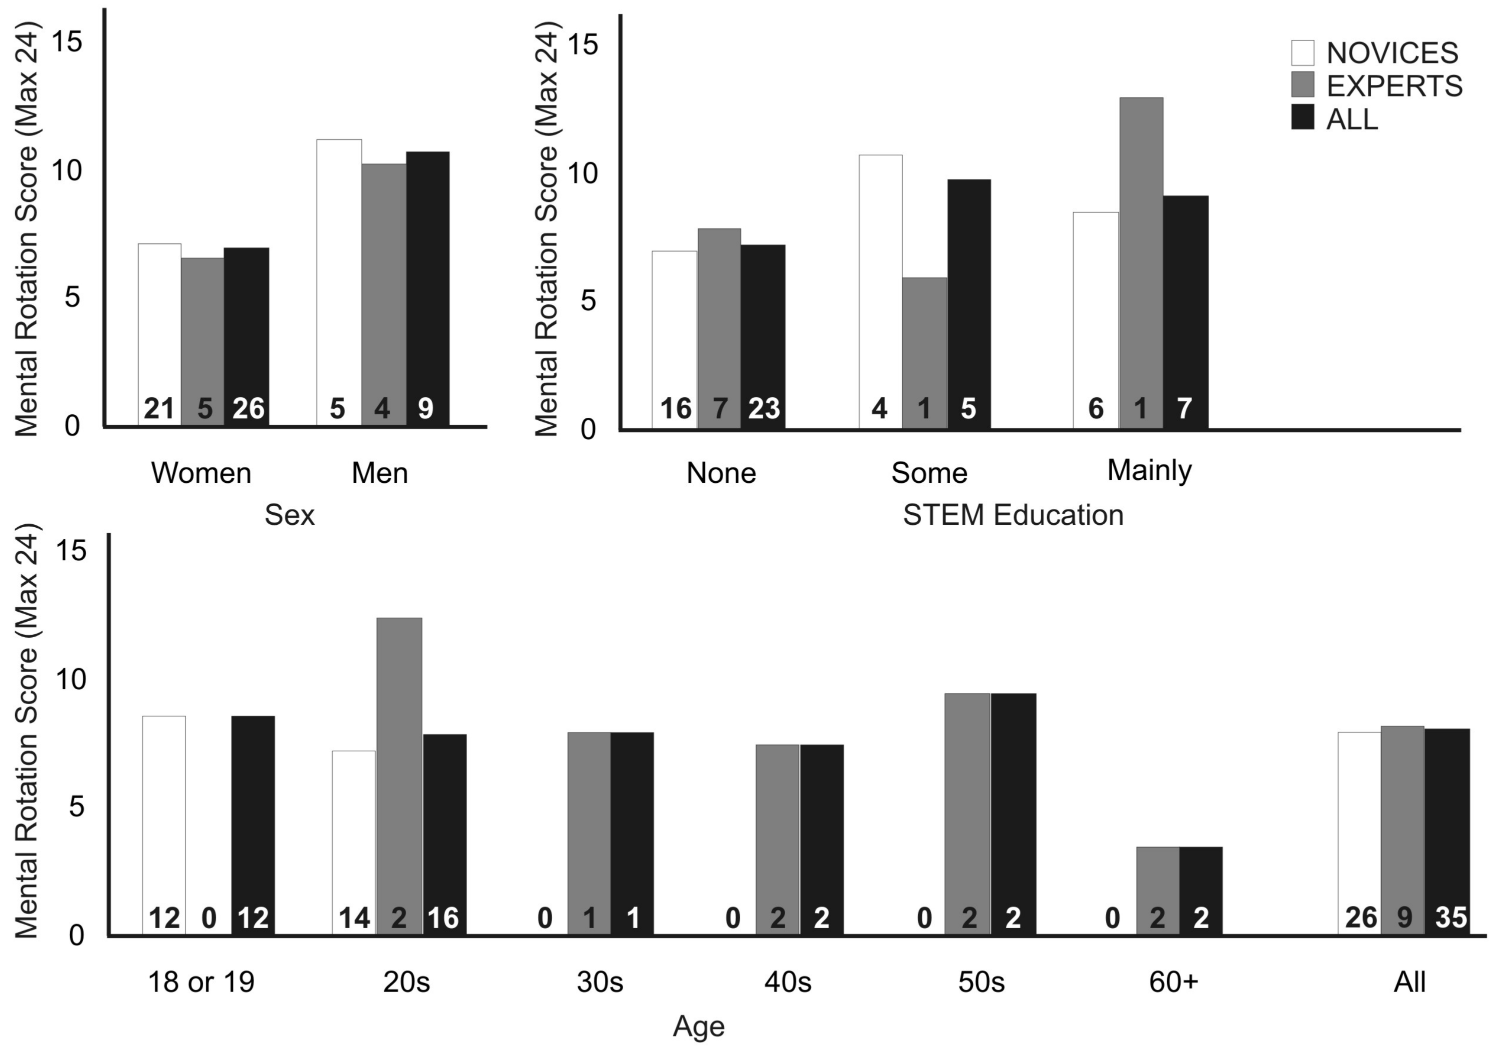

Figure C.1: There were uneven numbers of experts (N= 9) and novices (N= 26), and they were not matched on sex, STEM education or age. This may have confounded any pattern in the mental rotation scores. This figure reports the mean mental rotation score, by level of expertise within sex, level of STEM education and age. The numbers superimposed on each bar report the number of participants in that category. For example, there were 21 novice women and their mean mental rotation score was 7.14/24.

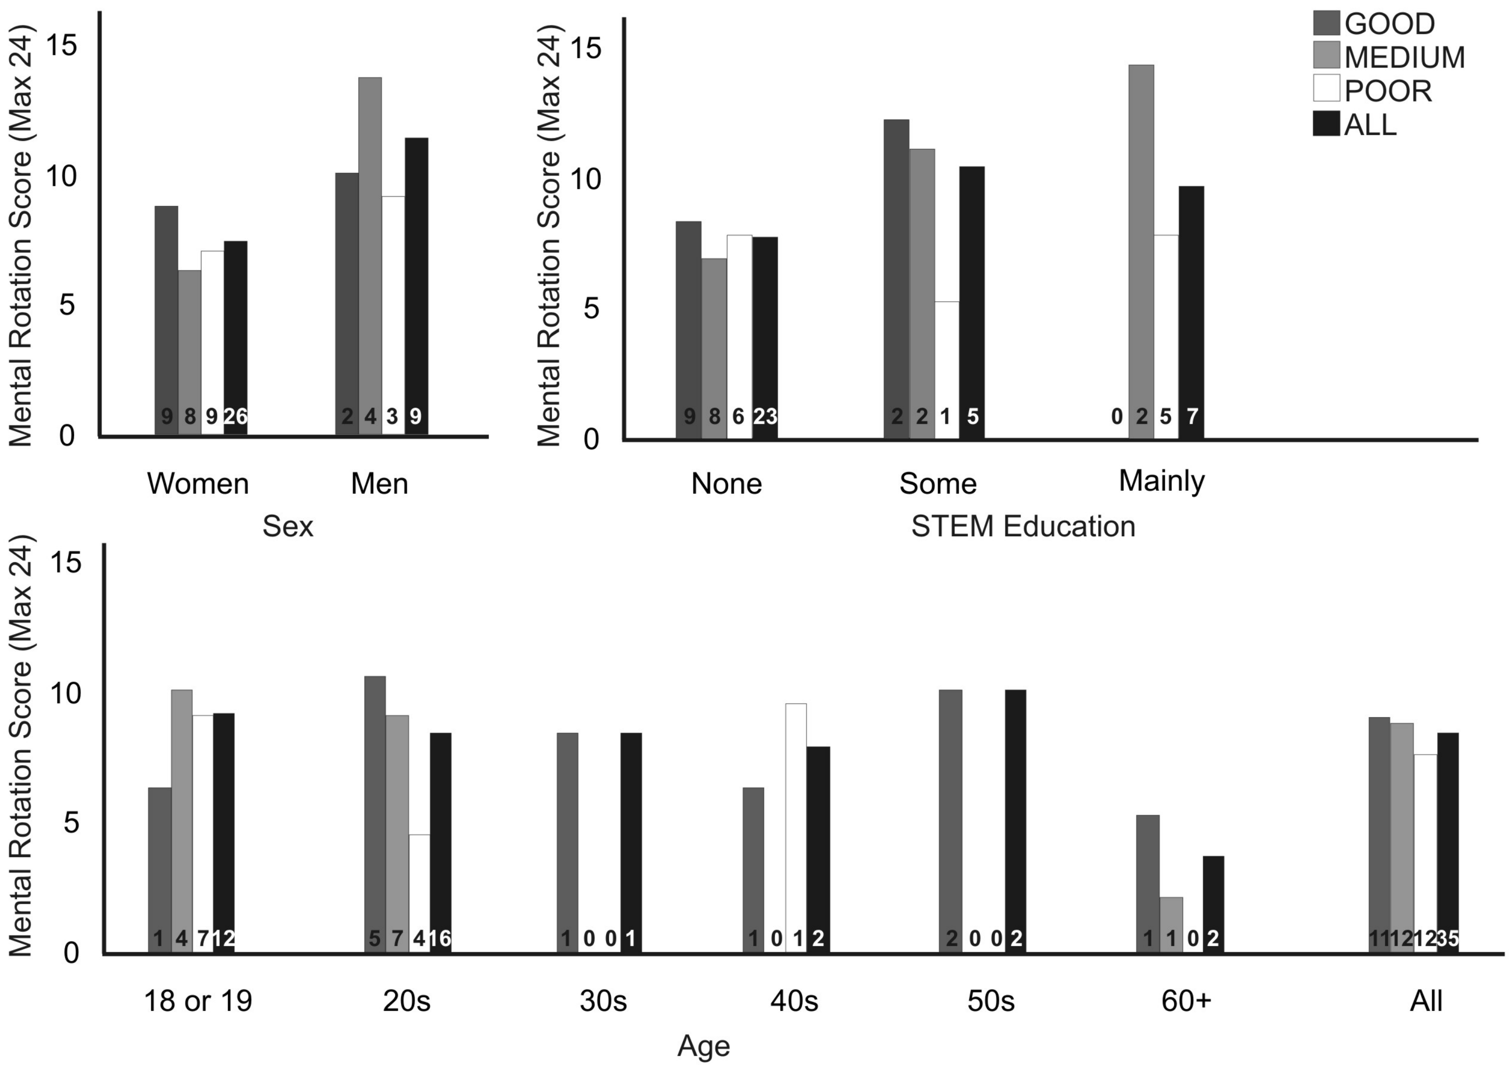

Figure C.2: When we stratified participants on accuracy into three even group sizes, membership was not matched on sex, STEM education or age. This may have confounded any pattern in the mental rotation scores. This figure reports the mean mental rotation score, by level of accuracy within sex, level of STEM education and age. The numbers superimposed on each bar report the number of participants in that category. For example, there were 9 women who drew with good accuracy, and their mean mental rotation score was 8.33/24.

Appendix D

Validating the choice of reference polygon

In the model in Chapter 3, most of the accuracy measures are calculated relative to a reference ROI which must be appropriately chosen. In order to confirm that the choice of Q_5 as the reference ROI had minimal effect on the model, we repeated the error measurements using Q_1 as the reference ROI. For each drawing, we compared the error measures relative to reference ROI Q_5 and reference ROI Q_1.

Because proportionality error is defined independent of the reference ROI, its values should be as close to identical as computation allows. As expected, the proportionality errors measured relative to Q_5 were identical to the proportionality errors measured relative to Q_1 for each of the three polygons.

The other three polygon-wise errors—orientation, scaling and position—are affected by the location, size and shape of the reference ROI. We should expect some minor variations with a change of reference ROI caused by the difference in the proportionality error of the reference ROI itself. For example, a change in the proportionality of the reference ROI will shift its centroid and minor principal axis, relative to which the polygon-wise errors are calculated. The drawing-wide means of these polygon-wise errors will also vary with the change in reference ROI because each mean includes the error from three common ROIs (Q_2, Q_3, and Q_3) and whichever of Q_1 and Q_5 is not the reference ROI. The mean polygon scale error of each drawing, when measured relative to reference ROI Q_5 and reference ROI Q_1, correlated significantly, $r(34) = .679$, $p < .001$. The mean polygon orientation error of each drawing, when measured relative to reference ROI Q_5 and reference ROI Q_1, correlated significantly, $r(34) = .656$, $p < .001$. The mean polygon position error of each drawing, when measured relative to reference ROI Q_5 and reference ROI Q_1, correlated significantly, $r(34) = .419$, $p < .05$.

When we compared the overall polygon error, $\mathbb{G}_i$ of each drawing, measured relative to reference ROI Q_5 and reference ROI Q_1, they correlated significantly with one another, $r(34) = .806$, $p < .001$. Having established that the model is satisfactorily independent of the choice of reference ROI, in Chapter 3 we reported the experimental results based on Q_5 as the reference ROI.

References

Adair, J., & Barrett, A. (2008). Spatial neglect: clinical and neuroscience review. *Annals of the New York Academy of Sciences*, *1142*(1), 21-43.

Allard, F., & Starkes, J. L. (1980). Perception in sport: Volleyball. *Journal of Sport Psychology*, *2*(1), 22-33.

Beery, K. E. (1968). Estimation of angles. *Perceptual and Motor Skills*, *26*(1), 11-14.

Bermingham, A. (2000). *Learning to draw: Studies in the cultural history of a polite and useful art*. New Haven and London: Yale University Press.

Bhattacharya, J., & Petsche, H. (2005). Drawing on mind's canvas: Differences in cortical integration patterns between artists and non-artists. *Human Brain Mapping*, *26*(1), 1–14. doi: 10.1002/hbm.20104

Biederman, I. (1987). Recognition-by-components: a theory of human image understanding. *Psychological Review*, *94*(2), 115-147.

Cain, T. I. (1943). The objective measurement of accuracy in drawings. *The American Journal of Psychology*, *56*(1), 32-53.

Carson, L. C., & Allard, F. (2008). Artists drawing angles: An expertise approach. *Journal of Vision*, *8*(6), 733.

Carson, L. C., & Allard, F. (in press). *Angle-drawing accuracy as an objective, performance-based measure of drawing expertise.*

Carson, L. C., Quehl, N., Aliev, I., & Danckert, J. (2011). Spatial visualization predicts drawing accuracy. *Perception*, *40*, 218. (ECVP Abstract Supplement)

Caxton, W. (1474). *Game and playe of the chesse.* Retrieved from http://www.gutenberg.org/ebooks/10672

Chan, D. W., Chan, L.-k., & Chau, A. (2009). Judging drawing abilities of Hong Kong Chinese gifted students: Could nonexperts make expert-like judgments? *Gifted Child Quarterly*, *53*(1), 15-24. doi: 10.1177/0016986208326555

Chatterjee, A. (2006). The neuropsychology of visual art: Conferring capacity. *International review of neurobiology*, *74*, 39–49.

Chen, S., & Levi, D. (1996). Angle judgment: Is the whole the sum of its parts? *Vision research*, *36*(12), 1721-1735.

Chi, M. T. H., Glaser, R., & Rees, E. (1982). Expertise in problem solving. In R. J. Sternberg (Ed.), *Advances in the psychology of human intelligence (vol. 1)* (p. 7-75). Hillsdale, N.J.: Erlbaum.

Cohen, D. J. (2005). Look little, look often: The influence of gaze frequency on drawing accuracy. *Attention, Perception, and Psychophysics*, *67*, 997-1009. doi: 10.3758/BF03193626

Cohen, D. J., & Bennett, S. (1997). Why can't most people draw what they see? *Journal of Experimental Psychology: Human Perception and Performance*, *23*(3), 609-621.

Cohen, D. J., & Earls, H. (2010). Inverting an image does not improve drawing accuracy. *Psychology of Aesthetics, Creativity, and the Arts*, *4*(3), 168-172.

Cohen, D. J., & Jones, H. . c. (2008). How shape constancy relates to drawing accuracy. *Psychology of Aesthetics, Creativity, and the Arts*, *2*(1), 8-19.

Cole, F., Golovinskiy, A., Limpaecher, A., Barros, H. S., Finkelstein, A., Funkhouser, T., & Rusinkiewicz, S. (2012, January). Where do people draw lines? *Commun. ACM*, *55*(1), 107–115. doi: 10.1145/2063176.2063202

Committeri, G., Galati, G., Paradis, A., Pizzamiglio, L., Berthoz, A., & LeBihan, D. (2004). Reference frames for spatial cognition: different brain areas are involved in viewer-, object-, and landmark-centered judgments about object location. *Journal of Cognitive Neuroscience*, *16*(9), 1517-1535.

Crawford, J., & Howell, D. (1998). Comparing an individual's test score against norms derived from small samples. *The Clinical Neuropsychologist*, *12*(4), 482-486.

Csikszentmihalyi, M., & Getzels, J. W. (1971). Discovery-oriented behavior and the originality of creative products: A study with artists. *Journal of personality and social psychology*, *19*(1), 47-52.

Dror, I. E., & Kosslyn, S. M. (1994). Mental imagery and aging. *Psychology and aging*, *9*(1), 90-102.

Dürer, A., & Strauss, W. L. (1977). *The painter's manual : a manual of measurement of lines, areas, and solids by means of compass and ruler assembled by albrecht durer for the use of all lovers of art with appropriate illustrations arranged to be printed in the year mdxxv / by albrecht durer ; translated and with a commentary by walter l. strauss.* Abaris Books, New York. (Original work published 1525)

Edwards, B. (1979). *The New Drawing on the Right Side of the Brain.* Los Angeles: Tarcher.

Ericsson, K. A., & Charness, N. (1994). Expert performance: Its structure and acquisition. *American Psychologist*, *49*(8), 725-747.

Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, *100*(3), 363 - 406.

Fayena-Tawil, F., Kozbelt, A., & Sitaras, L. (2011). Think global, act local: A protocol analysis comparison of artists' and nonartists' cognitions, metacognitions, and evaluations while drawing. *Psychology of Aesthetics, Creativity, and the Arts*, *5*(2), 135-145.

Ferrante, D., Gerbino, W., & Rock, I. (1995). Retinal vs. environmental orientation in the perception of the right angle. *Acta psychologica*, *88*(1), 25-32.

Fierro, B., Brighina, F., Oliveri, M., Piazza, A., Bua, V., Buffa, D., & Bisiach, E. (2000).

Contralateral neglect induced by right posterior parietal rtms in healthy subjects. *Neuroreport*, *11*(7), 1519.

Gauthier, I., Skudlarski, P., Gore, J. C., & Anderson, A. W. (2000). Expertise for cars and birds recruits brain areas involved in face recognition. *Nature neuroscience*, *3*(2), 191-197.

Gobet, F., & Simon, H. A. (1996). Recall of rapidly presented random chess positions is a function of skill. *Psychonomic bulletin & review*, *3*(2), 159-163.

Goldstein, C. (1996). *Teaching art: Academies and schools from vasari to albers.* Cambridge: Cambridge University Press.

Gombrich, E. (1960/1972). *Art and illusion: A study in the psychology of pictorial representation.* New Jersey: Princeton University Press.

Goodale, M., & Milner, A. (1992). Separate visual pathways for perception and action. *Trends in neurosciences*, *15*(1), 20-25.

Gray, R., & Regan, D. (1996). Accuracy of reproducing angles: is a right angle special? *Perception*, *25*, 531-542.

Gregory, R. L. (1963). Distortion of visual space as inappropriate constancy scaling. *Nature*, *199*(4894), 678-680. doi: 10.1038/199678a0

Halligan, P., Fink, G., Marshall, J., & Vallar, G. (2003). Spatial cognition: evidence from visual neglect. *Trends in cognitive sciences*, *7*(3), 125-133.

Halligan, P., & Marshall, J. (1991). Left neglect for near but not far space in man. *Nature*, *350*(6318), 498-500.

Halpern, D. F., & Wai, J. (2007). The world of competitive scrabble: Novice and expert differences in visuopatial and verbal abilities. *Journal of Experimental Psychology: Applied*, *13*(2), 79-94.

Harris, I., Benito, C., Ruzzoli, M., & Miniussi, C. (2008). Effects of right parietal transcranial magnetic stimulation on object identification and orientation judgments. *Journal of cognitive neuroscience*, *20*(5), 916–926.

Howard, I. P., & Allison, R. S. (2011). Drawing with divergent perspective, ancient and modern. *Perception*, *40*(9), 1017-1033.

Intraub, H., & Richardson, M. (1989). Wide-angle memories of close-up scenes. *Journal of Experimental Psychology: Learning, Memory, and Cognition; Journal of Experimental Psychology: Learning, Memory, and Cognition*, *15*(2), 179.

Jansen, P., & Heil, M. (2010). Gender differences in mental rotation across adulthood. *Experimental aging research*, *36*(1), 94-104. doi: 10.1080/03610730903422762

Jastrow, J. (1892). Studies from the university of wisconsin: On the judgment of angles and positions of lines. *The American Journal of Psychology*, *5*(2), 214-248.

Kottlow, M., Praeg, E., Luethy, C., & Jancke, L. (2011). Artists advance: Decreased upper alpha power while drawing in artists compared with non-artists. *Brain Topography*, *23*, 392-402. doi: 10.1007/s10548-010-0163-9

Kozbelt, A. (2001). Artists as experts in visual cognition. *Visual Cognition*, *8*(6), 705-723. doi: 10.1080/13506280042000090

Kozbelt, A. (2006). Dynamic evaluation of matisse's 1935 large reclining nude. *Empirical Studies of the Arts*, *24*(2), 119-137.

Kozbelt, A., & Seeley, W. (2007). Integrating art historical, psychological, and neuroscientific explanations of artists advantages in drawing and perception. *Psychology of Aesthetics, Creativity, and the Arts*, *1*(2), 80-90. doi: 10.1037/1931-3896.1.2.80

Kozbelt, A., & Serafin, J. (2009). Dynamic evaluation of high- and low-creativity drawings by artist and nonartist raters. *Creativity Research Journal*, *21*(4), 349-360.

Kubovy, M. (1986). *The psychology of perspective and renaissance art*. Cambridge University Press. Retrieved from http://people.virginia.edu/~mk9y/arrow/ppra.pdf

Lehmann, A. C., & Ericsson, K. A. (1996). Performance without preparation: Structure and acquisition of expert sight-reading and accompanying performance. *Psychomusicology: A Journal of Research in Music Cognition*, *15*(1-2), 1-29.

Linn, M. C., & Petersen, A. C. (1985). Emergence and characterization of sex differences in spatial ability: A meta-analysis. *Child Development*, *56*(6), pp. 1479-1498.

Lockard, W. (1982). *Design drawing*. New York: Van Nostrand Reinhold Co.

Logothetis, N., & Sheinberg, D. (1996). Visual object recognition. *Annual review of neuroscience*, *19*(1), 577-621.

Maclean, I., & Stacey, B. (1971). Judgment of angle size: An experimental appraisal. *Attention, Perception, & Psychophysics*, *9*(6), 499-504.

MacRae, A., & Loh, H. (1981). Constant errors occur in matched reproduction of angles even when likely biases are eliminated. *Attention, Perception, & Psychophysics*, *30*(4), 341-346.

Makuuchi, M., Kaminaga, T., & Sugishita, M. (2003). Both parietal lobes are involved in drawing: a functional mri study and implications for constructional apraxia. *Cognitive brain research*, *16*(3), 338–347.

Mann, D. L., Abernethy, B., & Farrow, D. (2010). Action specificity increases anticipatory performance and the expert advantage in natural interceptive tasks. *Acta Psychologica*, *135*(1), 17-23.

Matthews, W. J., & Adams, A. (2008). Another reason why adults find it hard to draw accurately. *Perception*, *37*, 628-630.

McManus, I. C., Chamberlain, R., Loo, P.-W., Rankin, Q., Riley, H., & Brunswick, N. (2010). Art students who cannot draw: Exploring the relations between drawing ability, visual memory, accuracy of copying, and dyslexia. *Psychology of aesthetics, creativity, and the arts*, *4*(1), 18-30.

McManus, I. C., Loo, P.-W., Chamberlain, R., Riley, H., & Brunswick, N. (2011). Does shape constancy relate to drawing ability? two failures to replicate. *Empirical Studies of the Arts*, *29*(2), 191-208.

Miall, R. C., & Tchalenko, J. (2001). A painter's eye movements: A study of eye and hand movement during portrait drawing. *Leonardo*, *34*(1), 35-40.

Milner, A., & Goodale, M. (2006). *The visual brain in action.* Oxford University Press, USA.

Mishkin, M., Ungerleider, L., & Macko, K. (1983). Object vision and spatial vision: Two cortical pathways. *Trends in neurosciences*, *6*, 414-417.

Mitchell, P., Ropar, D., Ackroyd, K., & Rajendran, G. (2005). How perception impacts on drawings. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(5), 996 - 1003. doi: 10.1037/0096-1523.31.5.996

Nicholls, A. L., & Kennedy, J. M. (1995). Foreshortening in cube drawings by children and adults. *Perception*, *24*, 1443-1456.

Nicolaides, K. (1941). *The natural way to draw: a working plan for art study.* Boston: Houghton Mifflon.

Ostrofsky, J., Kozbelt, A., & Seidel, A. (2012). Perceptual constancies and visual selection as predictors of realistic drawing skill. *Psychology of Aesthetics, Creativity, and the Arts*, *6*(2), 124-136.

Perdreau, F., & Cavanagh, P. (2011). Do artists see their retinas? *Frontiers in Human Neuroscience*, *5*(171). doi: 10.3389/fnhum.2011.00171

Perez-Fabello, M. J., & Campos, A. (2007). Influence of training in artistic skills on mental imaging capacity. *Creativity Research Journal*, *19*(2-3), 227-232.

Peters, M., Laeng, B., Latham, K., Jackson, M., Zaiyouna, R., & Richardson, C. (1995). A redrawn vandenberg & kuse mental rotations test: different versions and factors that affect performance. *Brain and Cognition*, *28*, 39-58.

Picard, D., & Durand, K. (2005). Are young children's drawings canonically biased? *Journal of Experimental Child Psychology*, *90*, 48-64.

Pratt, M. B. (1926). The visual estimation of angles. *Journal of Experimental Psychology*, *9*(2), 132.

Previc, F. (1998). The neuropsychology of 3-d space. *Psychological Bulletin*, *124*(2), 123-164.

Previc, F., & Intraub, H. (1997). Vertical biases in scene memory. *Neuropsychologia*, *35*(12), 1513-1517.

Rock, I., Wheeler, D., & Tudor, L. (1989). Can we imagine how objects look from other viewpoints? *Cognitive Psychology*, *21*(2), 185-210.

Roth, S. B. (1992). An investigation into cognitive factors involved in the drawing process. *Journal of Visual Literacy*, *11*(2), 57-76.

Sack, A., Camprodon, J., Pascual-Leone, A., & Goebel, R. (2005). The dynamics of interhemispheric compensatory processes in mental imagery. *Science*, *308*(5722), 702–704.

Sanjram, P. K., & Khan, A. (2011). Attention, polychronicity, and expertise in prospective memory performance: Programmers vulnerability to habit intrusion error in multitasking. *International Journal of Human-Computer Studies*, *69*(6), 428-439.

Schindler, I., Rice, N., McIntosh, R., Rossetti, Y., Vighetto, A., & Milner, A. (2004). Automatic avoidance of obstacles is a dorsal stream function: Evidence from optic ataxia. *Nature Neuroscience*, *7*(7), 779-784.

Schmidt, R., Khan, A., Kurtenbach, G., & Singh, K. (2009). On expert performance in 3d curve-drawing tasks. In *Proceedings of the 6th eurographics symposium on sketch-based interfaces and modeling* (pp. 133–140). New York, NY, USA: ACM. doi: 10.1145/1572741.1572765

Shepard, R. N. (1990). *Mind sights: Original visual illusions, ambiguities, and other anomalies.* New York: W.H. Freeman and Company.

Shepard, R. N., & Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, *171*(3972), 701-703.

Simon, A., & Ward, L. (1979). Some variables influencing judgments of the size of angles. *Contemporary Educational Psychology*, *4*(1), 11-14.

Snippe, H., & Koenderink, J. (1994). Discrimination of geometric angle in the frontoparallel plane. *Spatial vision*, *8*(3), 309-328.

Snyder, A. (2009). Explaining and inducing savant skills: privileged access to lower level, less-processed information. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *364*(1522), 1399–1405.

Snyder, A., Mulcahy, E., Taylor, J., Mitchell, D., Sachdev, P., & Gandevia, S. (2003). Savant-like skills exposed in normal people by suppressing the left fronto-temporal lobe. *Journal of integrative neuroscience*, *2*, 149-158.

Solso, R. (2001). Brain activities in a skilled versus a novice artist: An fmri study. *Leonardo*, *34*(1), 31-34.

Taylor, L. M., & Mitchell, P. (1997). Judgments of apparent shape contaminated by knowledge of reality: Viewing circles obliquely. *British Journal of Psychology*, *88*, 653-670.

Tchalenko, J. (2007). Eye movements in drawing simple lines. *Perception*, *36*(8), 1152-1167.

Tchalenko, J. (2009). Segmentation and accuracy in copying and drawing: Experts and beginners. *Vision research*, *49*(8), 791-800.

Tchalenko, J., & Chris Miall, R. (2009). Eye hand strategies in copying complex lines. *Cortex*, *45*(3), 368-376. doi: 10.1016/j.cortex.2007.12.012.

Thouless, R. H. (1931). Phenomenal regression to the real object. i. *British Journal of Psychology. General Section*, *21*(4), 339-359.

Ullman, S. (1998). Three-dimensional object recognition based on the combination of views. *Cognition*, *67*(1-2), 21-44.

Uttal, D. H., & Cohen, C. A. (2012). Spatial thinking and stem education: When, why, and how? In B. H. Ross (Ed.), *The psychology of learning and motivation* (Vol. 57, p. 356-378). San Diego, CA, US: Academic Press.

VanLehn, K., & Van, d. S. (2009). Acquiring conceptual expertise from modeling: The case of elementary physics. In K. A. Ericsson (Ed.), *Development of professional expertise: Toward measurement of expert performance and design of optimal learning environments* (p. 356-378). New York, NY, US: Cambridge University Press, New York, NY.

van Sommers, P. (1984). *Drawing and cognition: Descriptive and experimental studies of graphic production processes*. New York, NY, US: Cambridge University Press, New York, NY.

Voyer, D., Voyer, S., & Bryden, M. P. (1995). Magnitude of sex differences in spatial abilities: A meta-analysis and consideration of critical variables. *Psychological bulletin*, *117*(2), 250-270.

Wassermann, E. (1998). Risk and safety of repetitive transcranial magnetic stimulation: report and suggested guidelines from the international workshop on the safety of repetitive transcranial magnetic stimulation, june 5-7, 1996. *Electroencephalography and Clinical Neurophysiology/Evoked Potentials Section*, *108*(1), 1–16.

Wei, G., & Luo, J. (2010). Sport experts motor imagery: Functional imaging of professional motor skills and simple motor skills. *Brain research*, *1341*, 52-62.

Wenderoth, P., OConnor, T., & Johnson, M. (1986). Expansion and contraction of outline and solid acute angles: Effects of angle magnitude, type of display, and the nature of the matching task. *Attention, Perception, & Psychophysics*, *39*(4), 261-266.

Winter, D. (2005). *Biomechanics and motor control of human movement* (3rd ed.). John Wiley and Sons, Hoboken, NJ.

Yue, Y., Song, W., Huo, S., & Wang, M. (2012). Study on the occurrence and neural bases of hemispatial neglect with different reference frames. *Archives of Physical Medicine and Rehabilitation*, *93*(1), 156-162.

Zitov, B., & Flusser, J. (2003). Image registration methods: a survey. *Image and Vision Computing*, *21*, 977–1000.