New Approaches to Visual Scale and Visual Shape¹

Paul Linton

So the title of this talk is “New Approaches to Visual Scale and Visual Shape”.

And the purpose of this talk is to highlight 4 concerns with the 3D vision literature:

First, on visual scale, discussed my experimental papers.

Second, and this is new for this meeting, on disparity processing.

Third, on cue integration, discussed in my book.

And then fourth, on neural implementation, discussed in my recent V1 paper.

And they lead me to argue for an alternative to triangulation-based accounts of human 3D vision.

But in a sense the key question to explore in discussion is: “are these concerns shared by others in the field?”

So starting with visual scale.

How do we see the size and distance of objects?

Well this is explored in my papers “Does vision extract absolute distance from vergence” and “Does vergence affect perceive size”.

And the concern here is that all of our sources of absolute distance information appear to be either:

Ineffective

Very limited in application

Or merely cognitive.

So starting with familiar size.

¹ Talk at Royal Society meeting “New Approaches to 3D Vision”, 3rd November 2021.

Following Walter Gogel and John Predebon, it's generally thought that familiar size has a merely cognitive effect on our size and distance judgements.

The ground plane is very limited in application. For instance, if you're watching this talk at a desk, it doesn't seem to apply.

Renner and colleagues reviewed the literature on motion parallax.

And conclude that "there is no empirical evidence that providing motion parallax improves distance perception in virtual environments."

That's a bit harsh, but it's fair to say the evidence isn't convincing.

Mon-Williams & Tresilian found that accommodation, the focus power of the eye, provides "no functionally useful metric distance information".

Whilst tilt-shift blur is virtually never encountered.

Vertical disparities are the difference in the height of points between the two eyes.

But to be effective, they have to take up 60, 70 degrees of the visual field.

Since we rarely encounter objects that take up that much of the visual field, this can't be how we see scale in everyday vision.

Which only leaves vergence: triangulating distance using eye rotation.

Now there's a lot of evidence of vergence as an effective distance cue.

For instance, Mon-Williams & Tresilian find a strong linear relationship between the vergence distance of a point of light, and the distance participants pointed to with a hidden hand.

But I have 3 concerns with these and similar results.

And they all have to do with the way the stimuli are presented.

Because you're sat in darkness, and then a point of light is suddenly presented, as close as 20cm. And this is going to introduce 3 confounding cues.

First, you're going to see the point of light as double.

Second, as you converge your eyes, you're going to see the point of light move on the retina.

And then third, and I think this is just as important, if you suddenly have to converge your eyes as close as 20cm, you're going to be consciously aware of your own eye movements.

But we are virtually never aware of our own eye movements in everyday viewing.

So the purpose of my first experiment was to see what happens when we gradually vary vergence between trials, rather than suddenly at the beginning of the trial, so that participants are unaware their eyes are moving.

In this experiment, the left eye could see the right dot, and the right eye could see the left dot, specifying vergence.

And participants had to point to its distance with a hidden hand.

And between trials I would gradually vary vergence.

And participants would point again.

And what I found is that the effect of vergence on distance judgements is dramatically reduced.

And when you look at the raw data, it becomes immediately apparent that vergence is providing no useful distance information.

And the same is true for size perception.

So whilst Sperandio et al. (2013) find a strong linear relationship between vergence and the perceived size of an after-image.

I test this by having participants watch a target getting bigger or smaller and simply asking them: "Did the target get bigger or smaller?"

But the same time their vergence changes from 50cm to 25cm.

And because the change in vergence is supposed to be indicating that the target is getting closer and smaller, participants should judge the target to get smaller even when there is no physical size change.

But that's not what I find.

Take this one participant. When there is no physical size change, they don't judge the target to be getting smaller. Instead, they're at chance.

And this is true for all the participants. So vergence is having no effect.

So what are the implications of these results?

Well I think we go back to our list of distance cues.

And we cross vergence off the list.

And we simply conclude that the visual system isn't trying to estimate absolute size and distance.

But this poses a dilemma. Because we know that increasing the separation between the eyes makes the world seem smaller. So what explains this effect?

Well traditionally this has been attributed to vergence and vertical disparities.

But I've given you good reasons for being skeptical of this.

But the only other thing we're changing are the horizontal disparities.

And this doesn't seem a good candidate because horizontal disparities are thought to just affect the perceived depth of an object, but not its distance.

Well, I think that's right. But I also think that we learn to develop what I call "a cognitive association between (a) vivid stereo depth and (b) closer distances (reflecting our experience of an environment where disparity falls off with distance²)."

For instance, we only experience objects and scenes with accentuated depth at close distances.

Whilst a largely flat scene made up of largely flat objects is indicative of far distances.

So what this account relies upon is:

First, learning the relationship between accentuated depth and natural scene statistics.

But second, it also relies on extracting depth from disparity before distance.

Which leads us on to our next section on disparity.

Because it's the opposite of how we typically think about disparity processing.

Where distance is used to provide depth.

Rather than, as I am arguing, depth being used to provide distance.

So on the traditional account we first extract distance from vergence.

And then depth from disparity is simple triangulation.

But I've argued that we don't know this distance.

But one thing you might say is that is that the visual system just assumes a fixed distance.

But I don't think it's that simple.

And I want to raise three difficult cases.

The first is this stereogram. Or indeed most stereograms from the Victorian era.

Because they consist of uncrossed disparities viewed with parallel eyes.

So there's no solution to the triangulation problem.

The only solution is to pretend the eyes aren't parallel.

But first, it's not clear that that is going to solve the problem.

And second, triangulation loses much of its appeal.

You're severing the link between the position of points in the world and how they're processed by the visual system.

So those are impossible horizontal disparities.

But the second case I want to discuss are impossible vertical disparities.

So these arise in stereoscopic displays when something has gone wrong.

For instance, someone tilting their head without updating the cameras.

And what's really interesting is there's no real solution in the literature.

So many people just say "look, this is outside the scope of my account."

To their credit Held & Banks come up with an account that predicts that "the perceived shape of the front or back surface should become curved".

But they recognise that: "We have been unable to observe the curvature effect".

So those are impossible horizontal disparities, and now impossible vertical disparities.

But I think there's a deeper problem, and this is my third point, which is that the visual system "doesn't delete".

So what do I mean by this?

Well the promise of triangulation, whether it's in computer vision, or human vision, is novel view synthesis.

So we saw with Ali Eslami's talk in day one, you take some views of the scene, and you infer what the new view will be.

And the assumption is that the human visual system is doing something similar.

So you take the view of the left eye. And you take the view of the right eye.

And you synthesise a novel view, midway between the eyes, from what's known as the "cyclopean viewpoint".

Taking a view from the left eye, and a view from the right eye, and inferring a new "cyclopean" viewpoint midway between the eyes.

But that's not what happens.

Because it appears the visual system "doesn't delete".

Instead, it keeps both the view from the left eye.

And the view from the right eye.

And attributes them both to the cyclopean viewpoint.

And you can test this by having one eye open, looking at an object, opening your other eye, and seeing if you lose any of the view from the first eye.

But this doesn't make any sense from a triangulation perspective.

Because from a triangulation perspective the visual system should be giving us this. But instead its giving us this.

So something weird is going on. And I think we're entitled to think quite naively about what the explanation might be.

For instance, what the visual system does would make a lot of sense if the visual system didn't know that the views from the two eyes came from different viewpoints.

And there's something quite attractive about this, because it frames binocular vision as cyclopean from the outset.

There's no need to try and solve the complex problem of novel view synthesis, since there's only one viewpoint.

But can it explain depth from disparity?

Well, let's forget vision for the moment.

Instead, imagine these are two metal objects.

And I ask you to arrange them so that the circles overlap, and the dots overlap.

The problem is that if you overlap the circles, the dots aren't going to overlap.

But the secret is to realise there's a 3D solution to this problem. If you pull the arms up to form a cone the dots can also overlap.

And perhaps what the visual system is doing is something similar.

It's faced with binocular rivalry because it's got two inconsistent retinal images for the same viewpoint.

But just letting them rival is a worst-case scenario.

So it first tries to eradicate the rivalry by finding a 3D solution to this 2D problem.

If that's right, then binocular depth would be a solution to a problem, 2D binocular rivalry, rather than a problem to be solved on its own terms.

Now the solution the visual system adopts need not be the geometrically correct one. It could be this. It could be this.

And there's some evidence it's linear rather than geometric.

But how can we test whether something like this is correct, versus a traditional triangulation account?

Well, whatever the solution the visual system adopts, the geometry of the resulting percept should remain the same, no matter the size of the stereogram.

By contrast, on a triangulation account, changing the size of the stereogram should also distort its geometry.

Now I haven't had an opportunity to test this formally.

But we can at least construct stereograms to see if this looks right.

So you take a stereogram of a cone.

Halve it in size.

Halve it in size again.

And if the geometry is preserved then you should be able to stack those stereograms one on top of each other without distortion.

This is hard to convey via Zoom, so I've posted this talk as a preprint, and there you can find these stereograms to try for yourself.

So this is the simple version. And then the more complicated version.

Similarly, you can stack the stereograms without distortion.

Suggesting that if you double the disparity, you double the depth.

And that's illustrated by this stereogram, again in the preprint.

So one objection is to say that this only tests stereograms around the fixation plane, where there isn't that much distortion to be found.

But there's two responses:

First, we only ever experience binocular fusion around the fixation plane.

And second, if there's a simple mechanism that achieves the same outcome as triangulation, why prefer triangulation.

So to summarise this section, triangulation faces a number of problems.

No absolute distance from vergence.

Impossible vertical disparities.

Impossible horizontal disparities.

And the fact it doesn't delete.

But it's a historical mistake that we had triangulation for 200 years before we had binocular disparity.

And we've now spent the last 200 years trying to shoe-horn binocular disparity back into the triangulation account.

So perhaps the time has come to explore alternatives.

And treat depth from disparity as a solution to a different problem.

Binocular rivalry.

So that concludes the argument on binocular disparity.

I just want to make three brief points before we turn to Cue Integration.

First, if the visual system doesn't have access to vergence eye movements for distance.

Then it seems unlikely that the visual system has access to version eye movements for direction, since they're essentially the same signal, but the sign is changed in one eye.

Second, if the visual system doesn't have access to eye movements, this calls into question the visual system's role in visual stability, which is all about compensating for eye movements.

Third, the same experimental paradigms that are used for vergence are used to suggest that vision "relies on multimodal signals" and that "visual consciousness is shaped by the body". So my work calls this multisensory literature into question as well.

Ok, but let's continue to think about 3D vision.

You might argue that visual system must be doing 3D vision proper because it integrates binocular disparity with motion, perspective, and shading into a single coherent 3D percept.

But the argument of my book is that there's no evidence this integration is operating at the level of visual perception.

All I can do is give you a flavour of the argument with a few illustrative examples.

But Chapter 2, which contains the discussion, is available on my website.

So one example is Todd & Norman. Participants had to set the depth of a hinge to 90° in 3 conditions:

First, motion viewed with one eye.

Second, static binocular disparity.

And third, a motion plus disparity display.

Participants rated motion viewed with one eye as having most depth.

Static binocular disparity as having the least depth.

Least depth in the static binocular disparity display.

And, as predicted by depth cue integration, combined motion and disparity as having an intermediate level of depth.

And those are the objective results.

But during the debrief, they asked participants to watch the motion plus disparity display, and close one eye.

This turned it back to the monocular motion display.

But all the observers reported a reduction in depth.

Even though they'd previously rated it as higher.

So what's going on here?

The problem is that we're running two concerns together.

Pictorial depth and perceived depth.

Imagine I ask you to judge the depth in this cube.

You'll say its depth equals its width.

Then I ask you to judge the depth of this cube.

And you say its depth equals half its width.

So the cube on the left has more depth than the cube on the right.

But now imagine the cube on the right is defined by disparity.

You're still going to say its depth is half width.

So objectively, the perspective cube on the left has more depth.

But we know that isn't true.

Because the depth of the perspective cube is merely pictorial depth.

Whilst the depth of the stereogram on the right is perceived depth.

And the argument of the book is that in experiments we consistently mistake objectively reported pictorial depth for perceived depth.

Explaining why you find these inconsistencies between objectively reported results, and actual visual experience.

And this doesn't just apply to perspective.

It applies to the kinetic depth effect and motion parallax as we saw in Todd & Norman.

And looming, where an object suddenly expands in size.

For instance, in González et al., motion in depth from looming was judged to be roughly equal in magnitude to motion in depth from disparity.

And again, those are the objective results reported in the literature.

But observers still reported massive relative depth between the disparity and looming stimulus.

And the title of the book refers to this essential distinction between pictorial depth, which I take to be merely cognitive, and the actual perception of depth from disparity.

So what's the solution? Well, in my book I argue:

First, we need to remove time intervals between stimuli, so we're really testing perception and not working memory. That's the lesson to take from Todd + Norman.

Second, we need to stop testing stimuli in isolation. That's the lesson to take from Gonzalez et al.

Because we typically ask participants to evaluate the slant of a cue conflict stimulus viewed in isolation.

So in this case people judge it to have a slant somewhere between 0 (its slant defined by disparity) and 30 degrees (its slant defined by texture).

But we rarely view objects in isolation in the real world.

So if cue integration really governs our everyday perception, it should be able to persist in the presence of other objects defined by different sets of cues.

For instance, a fronto-parallel bar defined by disparity.

But at least to me, once an object defined by disparity is introduced, it becomes immediately apparent the surface defined by texture has no actual depth relative to the fronto-parallel bar.

So the key point is that if cue integration does not persist in the presence of other objects then it cannot be a theory of everyday vision.

So the final question I want to turn to is neural implementation.

And this is the topic of my recent paper "V1 as an egocentric cognitive map".

So the question is: "Is 3D visual space processed in V1?"

And what's interesting is that there appears to be a real tension in the literature.

The concern here is that there's real tension in the literature.

So on the one hand, there's evidence of very rich spatial processing in V1 from locomotion, as we saw with Aman Saleem's talk.

There's also evidence of multisensory processing.

And even high-level cognitive processing.

So you've got all of this rich spatial processing.

And yet it's thought that the most basic element of 3D visual processing, depth from disparity, happens after V1.

So that's the key tension. V1 is high-level enough to be doing all of this rich cognitive processing, but on the other hand it's too low level to be doing depth from disparity.

But this suggestion rests on what we think depth from disparity requires.

The first reason to reject depth from disparity in V1 are anticorrelated stereograms, where the colour is reversed between the two eyes

Because they provide evidence of disparity processing in V1 without a depth percept.

But on my account, that's to be expected.

Because depth from disparity is a way of avoiding rivalry.

But fusing these points in depth isn't going to eradicate their rivalry because they're different colours.

So you're better off just letting them rival one another as binocular lustre, which is what the literature finds.

A second reason to reject depth from disparity in V1 is that V1 processes absolute rather than relative disparity.

As Parker notes, they're looking for an account that "eliminates the fovea as a reference point" and that "avoids reference to any retinal landmark."

But on my account that's exactly what we want to preserve.

Because binocular rivalry is defined relative to the fovea.

Since whatever is on the fovea of the two eyes is fused.

And whatever is off the fixation plane, as defined by the fovea, rivals.

To conclude, I think it's useful to draw these themes together with an analogy.

Josh Tenenbaum talks about the “graphics engine in our head”.

But then the question is what does the visual system need to render?

And how does it render it.

The traditional answer is that it needs to render:

1. Size and distance of objects.
2. 3D Shape of objects.
3. It also needs to update the camera position with changes in eye position.
Perhaps with some smoothing for visual stability.

And we can think of this as Marr’s 2.5D sketch.

4. But neural implementation is hard to understand.

By contrast, on my minimal 3D account.

The visual system doesn’t process the distance of objects.

Only the depth of objects, and we cognitively infer distance on that basis.

Second, the visual system doesn’t even seem to be processing 3D shape.

And instead, could simply be eradicating binocular rivalry.

Third, since on my account eye movements don’t affect our visual experience, the visual system also doesn’t update the “camera position” with eye movements.

With the value of this account being that it makes it much more tractable how 3D visual space could be implemented in the brain.

Since everything in this account is consistent with V1.

So this talk is now available as a preprint on PsyArXiv.

And please join the debate on twitter @LintonVision

All that's left is to say thank you to everyone who helped with the experimental work. Christopher Tyler, Simon Grant, Matteo Lisi, Josh Solomon, Salma Ahmad, and Chris Hull. Thank you.