# Do we represent peripersonal space?

Colin Klein
The Australian National University
colin.klein@anu.edu.au

Work in both animals and humans has demonstrates that the brain specifically tracks the space near the body—the so-called peripersonal space (PPS). These representations appear to be multimodal and expressed in body-centered coordinates. They also play an important role in defense of the body from threat, manual action within PPS, and the use of tools—the latter, notably, 'extending' PPS to encompass the tool itself. Yet different authors disagree about important aspects of these representations, including how many there are. I suggest that the questions about the nature and number of PPS representations cannot be separated from the question of the mathematical basis of the corresponding representational spaces. I distinguish cartographic from functional bases for representation, suggesting that the latter is both a plausible account and supports a single-representation view. I conclude with reflections on functional bases and what they show about representation in cognitive science.

## 1. Introduction

The space near our bodies is important. It is where we act. It buffers us from threats. We spend a lot of time trying to keep people out of this space, and letting themn into it is a sign of trust and intimacy.

There is now considerable evdence that the brain contains neurons dedicated to the representation of this *peripersonal space* (PPS). Graziano and colleagues (1994) showed that the precentral gyrus of macaques contained neurons with bimodal visual and tactile receptive fields. Similar neurons are found in the ventral intraparietal area. These neurons respond to both touch on a part of the body and to a visual region extending out from that body part. The visual receptive field is anchored to the relevant body part, and moves as the body part moves. Some of these neurons are also responsive to nearby sounds, regardless of absolute sound intensity or whether the stimulus is visible (Graziano, Hu, and Gross 1997; Graziano, Reiss, and Gross 1999). There are more, and more densely overlapping, receptive fields closer to the body (Graziano and Cooke 2006 fig 4).

Further, stimulation of the same neurons using biologically realistic parameters can evoke complex defensive behaviors (Graziano, Taylor, and Moore 2002; Graziano et al. 2002; Graziano 2006). Stimulation of a region containing neurons with face-centered bimodal receptive fields, for example, evoked a complex defensive response involving head motion and a guarding response with the hand and arm, all directed towards the portion of space to which those responded (Graziano, Taylor, and Moore 2002). Though the organization of the precentral gyrus seems to involve a variety of different action-types, the polysensory zone seems to be especially concerned with defensive motions (Graziano et al. 2002). Reviewing the first two sets of findings, Graziano and Cooke (2006, 846) thus suggested that "a major function of these cortical areas is to maintain a margin of safety around the body and to coordinate actions that defend the body surface."

Further research has deepened our appreciation of these neural systems (see Bufacchi and Iannetti (2018) for a recent review). Yet this research has also raised significant empirical and conceptual questions about the nature of the PPS representation. This debate arises in part because the concept of representation itself remains the subject of considerable philosophical debate (Shea 2018).

In what follows, I will use PPS to tackle some of these outstanding questions about representation—both about peripersonal space in particular and about representation in general. The strategy will be oblique. I will start with the problem that PPS must solve: transformation between different coordinate systems. I will then show how the computational demands on coordinate transformation differ in important ways depending on the bases used for the two representational spaces being linked. I argue that there are good reasons to think that PPS uses what I'll call a *functional* basis (rather than a *cartographic* one) not the least of which is that it solves certain outstanding empirical problems. Finally, I will use this result to loop back around to questions of representation more broadly.

## 2. Coordinate Transformation and the choice of basis

Brozzoli et. al. make a crucial observation about our actions in PPS:

> …in order to interact successfully with the objects in the surrounding of our body, it is necessary to represent the position of the target object relative to the observer's body or body parts. Given that our hands can move simultaneously with and independently from our eyes, the brain needs to integrate and constantly update information arising in an eye-centered reference frame with information about the current position of the hand relative to the body and to nearby potential target objects. The perihand space representation provides an effective mechanism to support such a fundamental function. (2014, 130)

In other words, the primary problem that needs to be solved to enable effective action in peripersonal space is that of *coordinate transformation* between different representational spaces. Theories of PPS which otherwise disagree converge on this as something that PPS representations must do.

While I will mostly be concerned with the representation *of* space, it is important to note that the notion of a 'representational space' is far more general. Any determinable property which has $n$ different determination dimensions can be represented as an $n$-dimensional space, with the particular value taken by that property corresponding to a point in that space (Funkhouser 2006). If the dimensions which span this space additionally represent degrees of similarity, then overall more similar instances will appear closer together in that space. (Things can be similar in a variety of different ways, of course, and these different ways correspond to spaces with different bases). So for example, neuroscientists regularly talk of color spaces in perception, of feature spaces for recognizing objects, and motor spaces for action. None of these are concerned with representing space *per se*.

The coordinate transformation problem has four related but distinct parts (McCloskey 2009; Grush 2007 fn3). First, what corresponds to the origin in one space may not coincide with the origin in a different space. This is the natural reading of 'an $X$-centered reference frame.' So for example, an eye-centered reference frame might have its origin at the eye, while a hand-centered reference frame would have it at the hand. Events at the same location would have different coordinates in each frame.

Transformation of origins is a relatively trivial problem when the coordinate systems stand in a stable relationship. The second, more serious problem noted by Brozzoli et al. stems from the fact that the different origins can also move relative to one another. If the hand moves, the mapping from eye-centered to hand-centered coordinate space must also shift. Further, calculation of the relationship between these spaces cannot (in general) be derived solely from information within the two coordinate spaces. Coordinate transformation thus requires integrating additional information, like proprioceptive inputs about the location of the hand relative to the body.

The third, and more pressing, problem is that different coordinate systems can also use different bases. Formally, a basis is a mathematical notion: a minimal set of linearly independent vectors which span a given representational space.[1] Informally, the basis of a representational system can be thought of as the axes which define the space are meant to represent—its coordinate system, if you like. The same set of properties can be represented in different but equivalent ways, each of which can be thought of as corresponding to a representational space with a different basis.

So, for example, a retinotopic representation might use a two-dimensional polar representation, representing stimuli in terms of angle and distance from the fovea. An amodal allocentric representation might use something like a cartesian coordinate system with the origin centered on the head. An arm-centered representation might represent a location in space in terms of a multi-dimensional space with each axis representing a joint angle. Each of these represents a perfectly valid way to map out points in space. However, translating between these different bases can require requires complex and *nonlinear* mappings.

It is nonlinearity that makes this third problem difficult. Optimization of linear functions is a well-studied problem in both mathematical and neural contexts. Nonlinear functions, in contrast, are comparatively difficult to learn. Linearity is thus simpler from a developmental and a theoretical point of view. Of course, the coordinate transformations *must* be nonlinear in some sense. But all things considered, the fewer the nonlinearities one needs, the better.

The idea of a motor space is especially relevant for thinking about PPS, since at least some of the point of PPS is to guide motor actions in order to perform successful defensive motions. In present terminology that gives rise to the fourth problem: however PPS is represented, there is also a *prima facie* coordinate transformation problem between PPS-space and motor space. Motor space is probably not spatially organized: the coordinate axes are best understood either as involving parameters like joint angles and muscle contractions or else as more complex combinations of basic actions (Graziano 2016). Whichever way PPS is represented, then, there is an additional mapping required to get to the representational space necessary for motor output.

---

[1] I also assume that any particular representation has a privileged basis that serves as its coordinate system, which is not required mathematically but would seem to be necessary for any concrete instantiation of thae representation. This can complicate empirical inference; see Goddard et al. (2018) for an extended discussion of these issues as they arise in single-cell recordings.

## 3. Two strategies for representing space

Mathematically speaking, the choice of basis is usually irrelevant; there are many ways to represent the same space, each of which are sufficient for the job. *Computationally* speaking, the situation is quite different. How we represent space affects the complexity fo the computations performed with that representation. I will focus on representations of spatially organized information, but the points here are very general.

As an illustration, suppose that we are tasked with coordinating cartographic data that's given using different map projections. Each projection can be thought of as using a different basis, and correspond to a different way fo projecting a globe to a map.

Moving directly from one map projection to another is often mathematically complex. Further, the transformation that takes one map projection to another is typically irrelevant to any other pair of projections. As illustrated in figure 1a we could simplify our task by translating the information in each map back to a point on the globe that the map is meant to represent. This strategy not only simplifies the math, it is flexible: we can work with any projection that has a well-defined forward and inverse mapping to the globe. Adding a new projection only requires figuring out the mapping to and from the globe, rather than to and from every other projection we care about. The intermediate representation is also fairly easy for us to understand and work with: there is a simple isomorphism between a globe and the world it represents.

Call this the strategy of building a *cartographic basis*. A cartographic basis has (roughly) as many dimensions as the actual space we wish to represent. Each dimension corresponds in a relatively straightforward way to a single spatial dimension in the real world.

Contrast cartographic bases with a less familiar, but widely used, strategy for representing spatially organized information. A two-dimensional Fourier transform takes spatially organized information and represents it in the frequency domain instead. Variants of this strategy are widely used; the JPEG standard for encoding pictures (ISO/IEC 10918-1:1994), for example, uses a Discrete Cosine Transform (DCT). As shown in figure 1b, the encoding step of a JPEG image represents each 8x8 block of pixels as a combination of weights on each of the 64 basis functions. The inverse of this process can be used to recover the original image.

Call this second strategy that of finding a *functional basis*. Spatial information is represented as points in a high-dimensional space, the bases of which correspond to nonlinear functions of the original input.[2] The basis functions themselves are chosen for their useful mathematical properties, rather than for any obvious correspondence to the world. The dimensionality of this space is determined by the number of basis functions needed to span the original space with a desired degree of accuracy, rather than by the dimensionality of the target itself.

---

[2] To ease exposition, I will sometimes treat the strategy as involving a proper function space and sometimes as involving a space in $\mathbb{R}$ into which some original set of data is projected, with each of the basis functions corresponding to one dimension of the target space.

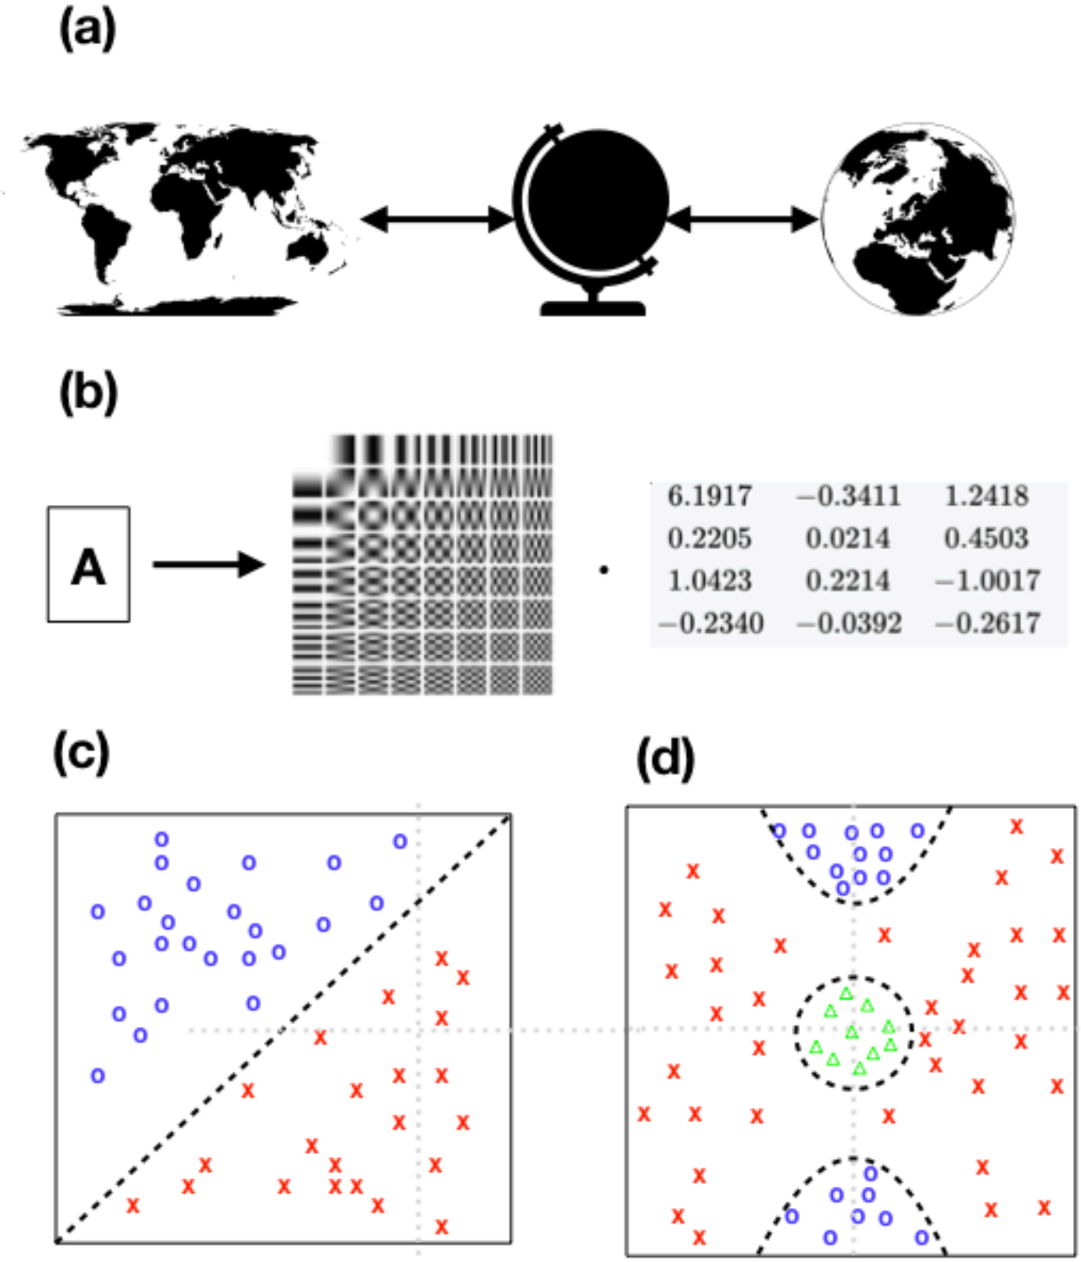

*Figure 1: (a) A globe represents a cartographic basis which simplifies movement between different map projections. (b) The discrete cosine transform (DCT)*[9] *representation of a simple letter. (c) A decision problem with a linear decision boundary. (d) Two decision problems with nonlinear decision boundaries: neither the noughts nor the triangles can be distinguished from the crosses with a linear cut.*

A functional basis is less intuitive and more complicated. Why go to the trouble? Because in a functional basis, certain algorithms that would be difficult to perform in a cartographic basis become trivial.

The functional basis for the JPEG standard images is chosen to facilitate compression: the high-frequency components of a picture can be dropped without much difference in image quality (at least to the human eye). Compression thus becomes a simple matter of bounding the representation. In a functional basis like the DCT, simple local operations on the representation also affect global properties of the output. Operations like noise removal, orientation detection, or smoothing can thus become relatively trivial.

Of course, there are tradeoffs involved: simple editing—say adding a circle around a point—can be surprisingly hard in a functional basis. The choice of basis thus involves tradeoffs which are determined by what you want to do with the resulting representation. There is always a close relationship between the way data is represented and the efficiency of algorithms performed upon it (Wirth 1976, 2).

A further advantageous application of functional bases is worth emphasizing. A large number of decision-making problems in machine learning involve finding a linear function that lets us distinguish different classes. Just as linear projections are preferable to nonlinear ones, linear separability is desirable for several reasons. Determining the best linear separator between points is a well-studied and relatively simple procedure: it can be done analytically in many cases, and is rapidly learned by a wide variety of neural networks. Nonlinear boundaries are more difficult to learn, especially in higher dimensions.

Sometimes the data is cooperative, as in figure 1c. That is, sometimes there is linear separation when we represent the data in a ‘natural’ basis where the dimensionality of the representation corresponds to the dimensionality of the problem space. This is often not the case. Figure 1d shows a case where the separating boundaries for two classes from a third are nonlinear. No straight line drawn through the space will separate the crosses and the noughts, even approximately.

One possibility for dealing with data like this would be to learn a nonlinear function. This nonlinear function, however, will be no help if we move to a different problem (say, distincguishing noughts from triangles in 1d). Whichever whatever problems nonlinearity introduces, it will recur with each new class. Instead, we can avail ourselves of a sneaky trick widely used in machine learning.

Note first that the separating boundaries in figure 1d are conic sections. Any conic section can be represented as $ax^2 + by^2 + cx + dy + e$. Suppose we project our inputs $x$ and $y$ to a new four-dimensional functional basis $W$, with axes corresponding to $\{x^2, y^2, x, y\}$. Any conic section in our original space corresponds to a hyperplane in $W$. Thus in $W$ the classes we care about become linearly separable. In other words, when faced with a nonlinear decision boundary to learn, we can sometimes push that nonlinearity into the space in which the problem is represented; then linear boundaries of the new space correspond to nonlinear boundaries in the original space.

The basis of $W$ were chosen solely for their mathematical properties. $W$ is larger than it needs to be to solve any individual classification problem. However, having picked a single, slightly more complex representation, we can then solve innumerable different categorization problems by simply learning linear functions in $W$.

Along the same lines, note that the class boundaries can be modified in useful ways simply by biasing the corresponding linear functions from $W$. Suppose, for example, we wanted to stretch the class boundary of the triangles in 1d so as to be more elliptical along the x-axis. We can keep $W$ fixed and simply change the weight corresponding to the $x^2$ axis. As this comes at the readout stage, not in $W$ itself, we could thus stretch the class boundary without affecting our projection from any other other class.

Functional bases are thus powerful tools for representing the world. By systematically projecting input to a higher dimensional space in a nonlinear way, an indefinite variety of decision problems can be represented as linear maps from that space. In practice, the dimensionality of the spaces of interest will be much higher, and the most useful bases will be specific to the set of problems at hand.

## 4. PPS is represented by a functional basis

In the opening section, I canvassed some reasons to think that there are neurons specialized for representing PPS at all. But that does not really settle the question of the basis of that representation. Most work on PPS, at least until recently, has tacitly assumed a cartographic basis. So, for example, Bufacchi and Iannetti (2018; 2019)'s recent presentation in terms of fields surrounding the body is most naturally read as asserting that the internal representation of PPS is structured along broadly cartographic lines, with action values specified at each spatial location.

I think there is a reasonably strong argument that PPS is represented using a functional basis. The argument is primarily from computational parsimony. Given the variety of things that PPS representations are invoked to do, a cartographic basis is forced to proliferate PPS representations and their interactions. Each of these new representations introduces new learning challenges. A functional basis, by contrast, allows for a fixed and fully general basis which can be used across a variety of contexts, and which faces only linear learning problems. Absent reasons to believe otherwise, I think this is decent evidence for a functional basis.

To begin, the neural evidence clearly suggests that there are nonlinear interactions at the representation stage. Anderson and colleagues noted that successful visually-directed action of external objects must involve combining information from three different coordinate frames: head, eye, and retina. They found neurons in the macaque parietal lobe that were best modelled as involving a multiplicative–i.e. nonlinear—interaction between retinal and head position (Andersen and Zipser 1988). Andersen et al. (1993) extended this work, showing similar nonlinear interactions with bodily position.

Anderson and colleagues took this as evidence of multiple, distributed spatial representations in the brain, corresponding to distinct coordinate systems. However, Pouget and Sejnowski (1997) showed that these data could be better accommodated by a model in which there is a single representation that uses nonlinear basis functions to represent the location of objects. The nonlinear transformations used are different from the ones considered in the previous section: products of gaussian functions of retinal position, and sigmoidal functions of retinal location (1997, 224). A suitable ensemble of these functions can combine and encode motor input in a way that is suitable for *linear* readout by the motor system. Further, in such a model a stimulus "is represented in multiple frames of reference simultaneously by the same neuronal pool" (p223). This explains, among other things, why hemineglect due to damage of this pool of neurons does not seem to be confined to a single frame of reference.

Of course, functional and cartographic bases must involve some nonlinearity: both require a nonlinear projection from earlier sensory areas to the representation of PPS. I take it that the significance of the neural evidence is twofold. First, it shows that there are neurons with the relevant sorts of properties to form a functional basis. Second, it shows that these neurons collectively form a basis suitable for overall functional representation.

Consider next the computations that would be needed to efficiently generate the actions for which PPS is posited in the first place. PPS is implicated in (among other things) tool use and defensive actions. PPS neuron receptive fields are plastic, and can be altered

A second set of evidence comes from the plasticity of PPS receptive fields. There is both neural and behavioral evidence that the zone of PPS is plastic and can be altered in response to both task and stimulus. Receptive fields in the ventral intraparietal area can be altered by spatial attention (Cook and Maunsell 2002). Bufacchi and Iannetti (2018) canvass a variety of ways in which defensive motions need to be context-sensitive. When faced with a hungry tiger, the decision to flee or to climb a tree depends crucially on the distance to both the tiger and to the nearest tree.

Yet this plasticity puts constraints on the underlying representation. De Vignemont and Iannetti (2015) note, tool use and defensive actions give rise to fundamentally different kinds of action, with different ends, different trajectories, and different informational grain. Defensive actions are fast and often automatic, whereas skilled actions are often slow and conscious. The modulation of PPS also differentially affects defensive and tool-using actions. While there is some evidence that tool use can extend the defensive peripersonal field (Rossetti et al. 2015), the two kinds of actions would seem to require fundamentally different sorts of plasticity. Indeed, as Povinellia, Reaux, and Frey (2010) note:

> …during most activities of daily living, tools and utensils are used to perform actions within our natural peripersonal space. Moreover, tools are frequently used in ways that we would never employ our hands. For instance, we will readily use a stick to stoke the hot embers of a campfire, or stir a pot of boiling soup with a wooden spoon. In these circumstances, the target of the actions may be located well within reach, but a tool is chosen as a substitute for the upper limb in order to avoid harm. These examples suggest that we maintain separate non-isomorphic representations of the hand vs. tool as concrete entities even when using handheld tools within our normal peripersonal space (243-4).

This suggests that the defensive PPS representation should remain unchanged even when a skilled PPS representation is extended by tools.

So far, so good. Now, either sort of basis can accommodate these behavioral facts. However, the malleability of PPS presents something of a challenge if one assumes that PPS is represented by a cartographic basis. On such accounts, the change in PPS is typically modeled as deriving from a bias or distortion of the underlying representation itself. Anxiety (say), warps the representation of PPS in subtle ways, with the result that we act as if there were larger boundaries around objects.

That distortions of PPS are effected by distortion of the underlying representation is, I suggest, the natural reading when using a cartographic basis. The input and output transformations are complex and nonlinear, and subtle variation of these mappings would be difficult to manage. Simply changing the space—representing a spider as bigger than it actually is, for example—is the only way to allow the hard-won nonlinear mappings to stay stable while putting linear biases on parts of the represented space itself.

De Vignemont and Iannetti (2015) thus draw the natural conclusion: if PPS is altered by altering the underlying representation, then there cannot be a *single* PPS representation (see similarly Bufacchi and Iannetti (2018)). The effect that different stimuli have on PPS appear to differ depending on whether we consider defensive or tool-use behaviors . This would create contradictory demands on modification of the representation of PPS. Hence, they argue, there must be at least two distinct representations of PPS, covering different use cases.

Yet proliferating PPS cannot stop there. There are a great number of circumstances in which the line between tool-use and self-protection breaks down. Many tools, especially when working close to the body, can injure. Flint knapping, for example, involves working with sharp instruments close to the body, and comes with well-attested possibilities of serious injury.[3] Success in such situations requires close coordination between defense and tool use. Similarly, we do actually protect tools in many cases: a butcher's knife will chip and break on bone, a poorly handled marshmallow stick will catch fire. This may be especially pressing with premodern, hand-crafted tools.

These problems are compounded when there is cooperative action with tools, such as in hunting: the polite hunter avoids spearing himself *and* his companions. Indeed, one should probably treat the standard cases—objects unexpectedly threatening the body, and uncomplicated tool use in constrained situations—as the limit cases of more realistic, complex situations. Finally, there are a number of situations that don't seem to fit neatly into the defensive/tool use dichotomy. Consider walking through crowds, or trying to punch someone nearby: sometimes the body is both the tool and the thing being protected.

The cartographic strategy, therefore, must proliferate representations and interactions in order to deal with the complexity of PPS behavior. Again, this is a direct computational consequence of the choice of basis. Since (*ex hypothesi*) the mapping from sensory organs to PPS-space and from PPS-space to motor space are both nonlinear, it is difficult to see how consistent alterations to either could easily be made. Conversely, since representations in a cartographic basis bear a relatively simple relationship to external space, it is easy to see how alterations of the PPS representation would fit the bill. This means that the only convenient point for simple alteration is at the representation of PPS itself—but then we must proliferate representations, since there are too many different ways in which PPS can be altered.

Contrast this with the use of a functional basis. Figure 2 shows the overall picture. Different coordinate frames are projected to a single, stable high-dimensional representation. This can then be projected linearly into a variety of different coordinate frames (or into the same frame in

---

[3] See Lycett, Cramon-Taubadel, and Eren (2016) for ethnographic references, and Whittaker (1994 Ch5) for some vivid modern examples.

different ways) by simple linear readouts of the functional space. A single representation can be used in a variety of simple ways. Further, by assumption, linear transformations are comparatively easy to learn, so the proliferation of *readouts* for different tasks presents no deep difficulties.

Similarly, linear readouts from a functional basis provide an easy method for biasing motor output via simple linear biases on the readout functions themselves. And since there can be can be multiple independent readout functions from the same functional spac, different demands can give rise to distinct biases on this representation. Indeed, though I propose the model to account for PPS activity, there's no particular reason why it need be confined to PPS. Any nonlinear function from the same inputs to the same output might be profitably represented as linear readouts of this space, so long as the nonlinearities involved are in the span of the basis functions themselves.

The overall picture is one in which the core representation of PPS is a high-dimensional space that takes as inputs nonlinear functions of head position, limb position, eye position, and so on. The nonlinear basis functions that underly that transformation are chosen so that the readout end—here, I'll assume the readout is done by various aspects of the motor system—is a linear function of the PPS space. These linear readouts can in turn be biased by different contextually-sensitive inputs. Figure 2 puts everything together.

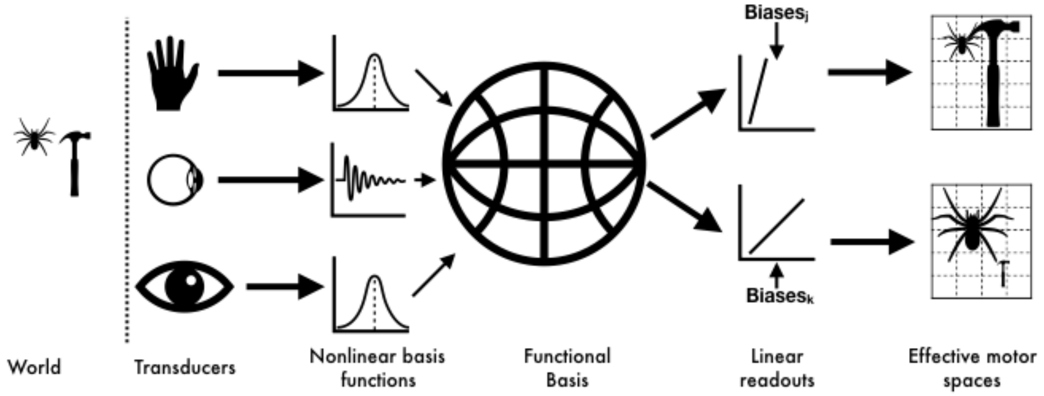

*Figure 2: A schematic picture of the functional basis model of PPS.*

To sum up the argument: the choice of basis in which PPS is represented matters because it makes certain kinds of computational operations easier or harder. There is considerable evidence that PPS is represented by neurons with the right sorts of properties to underly a functional basis. A cartographic basis has relatively restricted resources with which to accommodate shifts across different contexts. In addition, a cartographic basis would require the individual to learn two different, specific nonlinear mappings in developmental time. This adds additional complexity. By contrast a functional basis requires learning only a generic nonlinear mapping into functional space, and then a series of linear mappings from that space. Thus a functional basis ends up apparently simpler and more parsimonious across a variety of fronts. Of course, computational parsimony arguments for empirical findings are not conclusive, but all things being equal I suggest that parsimony provide a useful set of constraints on available theories.

## 5. Do we represent peripersonal space?

I conclude by returning to the question of representation. I have spoken so far as if it were obvious that we represent peripersonal space. Yet even if we keep the empirical evidence fixed, there remain several interesting questions which turn on broader points about the nature of representation itself. The question of whether the brain represents peripersonal space admits of three distinct readings, each with different emphasis, each important. I consider each in turn.

### 5.1 Peripersonal representation

First, we might want to know whether the brain really represents *peripersonal* space. That is a question about whether the representations we use are specially tuned to PPS itself. For all I've said above, this is not obvious. It might be that we simply represent space and threats and tools, and the space near our body happens to be a place where these interests coincide. The differential representation on the space close to the body might then just be a side effect of the fact that most relevant action categories happen nearby.

Indeed, this is a possibility which has been explicitly raised in the recent literature. In a recent review, Bufacchi and Iannetti (2018) suggest that PPS is represented by a series of 'fields' which (like their analogues in physics) are defined at every point in space. Noel and Serrino (2019) protest that this is no longer a theory of peripersonal space as such, because it does not confine itself to the space near the body. In response, Bufacchi and Iannetti (2019) appear to concede the point, arguing instead that PPS representations merely serve to protect the body, which might require tracking distant objects. But that would appear to stretch the definition of 'peripersonal space' to cover nearly any space. Conversely, Noel and Serrino (2019) 's argument appears to rely on the claim that certain action values are 'explicitly hard-coded in PPS neurons.' This strikes me as inconsistent with the previously canvassed arguments for behavioral and neural plasticity.

I think a focus on representational basis suggests a useful—and empirically testable—intermediate position. The question, it seems to me, turns on what PPS representations track. In a fully general case, where the functional basis covers all of space with equal fidelity, there would be a good argument that PPS is not represented as such. Conversely, if the functional basis covered *only* PPS—that is, if it only enabled activity in nearby space—then we would have good evidence that PPS as such is represented. We have good reason to think that the latter is not the case, because of the extension of PPS by reaching tools.

Instead, an intermediate possibility is likely to be instantiated. Pouget and Sejnowski (1997, 233) note that basis function representations with multiple input dimensions and uniform coverage can get large rather quickly. The power of functional bases comes from their increased dimensionality, but that high dimensionality comes at a cost as well. One possible solution they suggest is that the parietal cortex might "selectively span the input space to achieve greater efficiency."

There is evidence that this is what occurs with PPS. As Graziano and Cooke (2006) note, the space near the body appears to be covered by more, and more densely overlapping, receptive fields (see figure 4). Space further from the body is still represented by multimodal neurons, but with lower fidelity. Assuming this represents the organisation of the underlying functional basis, this might represent a compromise between full coverage and inflexibility. Conversely, this would also set limits on the degree to which further and further portions of the bodily space could be incorporated into PPS, and the accuracy with which the resulting incorporation might be achieved.

Here it strikes me that the decisive factor might involve experimentation with a wider variety of tools. Most experiments with tool use tend to focus on ethnologically realistic reaching tools that still operate near the body. Extremely long or large tools—telescoping tree pruners, or hydraulic excavators, for example—represent a novel empirical arena for testing the degree to which PPS can be extended far past the usual limits. The further PPS can be extended, the better an argument that there is no PPS *as such*, only regions of space where there are typically important actions to be done. Conversely, evidence of a spatially restricted area of high fidelity would suggest a specific and relatively inflexible tuning for near space, albeit one that can manage coarse and low-fidelity representation further afield.

### 5.2 Spatial representation

Following from the first question, we might ask a second, more general one. Supposing the brain represents something relevant to peripersonal defense (and the like), we might ask whether the brain represents peripersonal *space*. Our representation of PPS might be best understood as a representation of affordances for action in PPS (say) rather than as a representation of the space itself. Indeed, one thing that makes questions about PPS so interesting is that it is not at all clear that there is a distinctive phenomenology of PPS. That is, what is most salient in PPS (at last when I introspect) is usually not space itself but objects, tools, threats, and the like.

Focus on representational bases lets us reformulate this question in a useful way. There is a clear sense in which a cartographic basis represents space. But it is not obvious whether functional basis can be said to represent space as such. The constraints on a functional basis are primarily that of mathematical convenience, rather than similarity to external space. What's more, as one of the goals of a functional basis is to allow for simple linear representations of output, one might claim that the point of a functional basis is to represent stimuli *in terms of* behaviors that might be performed on them.

Grush, for example, reads Pouget's model as having this sort of structure (Grush 2007; Grush 2009). He claims that "The basis functions' entire purpose is to support a linear decoding via a motor-behavior- specific set of coefficients to produce a bodily action that is directed on that stimulus—such as grasping or foveating the seen object"; the objects in question "are objects in the sense of a potential focus of perception and action," and the functional basis as a whole underlies the "capacity to represent an environment of actionable objects" (Grush 2009, 342–43).

While I can see the attraction of such a view, I think it is mistaken. It is true that PPS is very often used for detecting and acting upon affordances. So a functional basis *plus* a readout function might be considered as something like a representation of affordances. But, as I've set it up, a functional representation of PPS is also a *decoupled* representation in the sense of Sterelny (2003). It can be used for a variety of different readout functions in different contexts. Further, since it forms a complete representational space, it is fully general. Together, this means that the representation of PPS is useful for an indefinite number of tasks. And what seems common to all of them is the concern with spatial location. Thus even though there's no obvious isomorphism between a PPS representation with a functional basis and space itself, the use of that representation means that it is best understood as spatial.

### 5.3 Representation as such

Third and finally, we might ask whether the brain *represents* peripersonal space. That is both a philosophical question about the nature of representation and an empirical question about how the brain tracks PPS. What the brain does with peripersonal space might fall short of fully representational: it might be, for example, that our facility with peripersonal space is best considered as a species of online skill, and that online skill is itself nonrepresentational. This is the sort of view that might be championed by non-representationalist strands in philosophy of mind (Chemero 2009). I think this is too hasty, and ultimately depends on a false dichotomy. Again, focus on the particular properties of functional bases is helpful.

The use of a functional basis for PPS was motivated in part by the problem of coordinate transformation. Yet coordinate transformation itself doesn't seem like it necessarily involves representation as such. Broadly speaking, it seems like coordinate transformation is more like a computational function than a representation per se. Webb and Wystrach (2016, 29), discussing place learning in cockroaches and crickets, contrasts learning strategies that are sufficient to do simple coordinate transformations yet that fall short of "internal representation of the spatial layout" of an arena.[4] Obrien and Opie (2009) distinguish digital representations, which stand in a completely arbitrary relationship to their target, and analog representations, which stand in a resemblance relationship to their target. Functional bases don't seem like they fit either bill, which again suggests that representation might be an inappropriate category.

Yet this argument relies on an overly restrictive notion of representation as something that is wholly explicit and always present to stand in for the original stimulus, whether or not that is being used at the time. But the primary function of a representation in more complex tasks is always transformative; the goal is to take information and re-present it in a form that is more amenable to further use (Ritchie 2015). DiCarlo and colleagues (2007; 2012), for example, model visual object recognition as a series of progressive transformations of information in the retinal image in a way that makes latent information simply accessible.

As before, I suggest that the more basic feature of representations is that they are decoupled from the immediate stimulus (Sterelny 2003) but still provide the same systematic handle on the world that the original stimulus could provide. Systematicity is key. Mere mappings from a single stimulus to a single response aren't representational, even if the mapping is complicated. It is

---

[4] More precisely, the contrast is with the strategy of learning a single exemplar of a panoramic view and then comparing it with the present view via gradient descent. For a stronger statement, see also Webb (2006). I think one can argue that even this actually tacitly represents a space (compare to the so-called 'kernel trick' in machine learning that allows computation of a gradient in space without requiring an intermediate transformation to that space), but that's stronger than needed for the moment.

systematicity that allows the same representation to be used for a variety of different responses in a variety of circumstances.[5]

It is clear that both cartographic and functional bases are decouplable. Both are equally concerned with combining information from the input end into a form that can be used in a variety of different ways. Indeed, the entire motivation behind moving to a functional basis was to facilitate decoupled use in distinct contexts. Of course, this does not settle the philosophical debate about the fundamental nature of representations are, but it suggests that both sorts of bases are similar enough that we ought to feel comfortable treating both or neither as representations.

That said, I think an assumption that PPS has a cartographic basis tacitly suggests a picture on which PPS is *explicitly* represented: that is, that everything about local space must be present like a map somewhere in cortex. There is no doubt that explicit maps are useful for some cognitive tasks (Rescorla 2009). But as Bufacchi and Iannetti (2019) point out in the defense of their field theory, what's really important is being able to systematically recover and transform action values for specific points in space. Describing a space with a functional basis also makes it more clear that explicit representation is less important than decouplable, systematic representation for most tasks.

It strikes me that the focus on PPS and coordinate transformation would be a good place to help with this, as there are a great number of cognitive tasks which can plausibly be modeled as species of complex coordinate transformations. There is a well-known body of work on the more general issue of egocentric to allocentric transformations of space, of course (Klatzky 1998). Effective action requires intertranslation between representations of outcomes and representation of motor actions (Butterfill and Sinigaglia 2014). Even some apparently higher-level judgments, such as sterotype-congruency, might be usefully modeled in terms of biases on underlying transformations, rather than explicit representation of stereotypes as such.[6] Focus on PPS, and on the possibility of functional bases for its representation, opens up a host of options that go beyond a simplistic distinction between transformation and representation.

# Conclusion

I have argued that we represent peripersonal space, but that the structure of that representation is best understood as involving a functional rather than a cartographic basis. We represent PPS in the sense that we have an efficient means for coordinate transformation that allows for the efficient learning and execution of PPS-related behaviors.

I take it that there is a broader lesson to be drawn. Debates over whether we represent PPS are ultimately questions about the number and taxonomy of processes involved in dealing with our behavior near our bodies. That is, they are questions about the *cognitive ontology* that we need to

---

[5] For a full presentation of this argument and its ramifications, see Klein and Clutton "Representation versus Transformation in Bodily Representation", in preparation.

[6] This argument is made forcefully by Gabbrielle Johnson, "The Structure of Bias", in review.

posit in order to make sense of observed behavior (Janssen, Klein, and Slors 2017). Thinking about PPS shows that questions about the number of processes often cannot be separated from the question of the format of the representations upon which the processes operate.[7]

## References

Andersen, Richard A, and David Zipser. 1988. "The Role of the Posterior Parietal Cortex in Coordinate Transformations for Visual–motor Integration." *Canadian Journal of Physiology and Pharmacology* 66 (4): 488–501.

Andersen, Richard A, Lawrence H Snyder, Chiang-Shan Li, and Brigitte Stricanne. 1993. "Coordinate Transformations in the Representation of Spatial Information." *Current Opinion in Neurobiology* 3 (2): 171–76.

Brozzoli, Claudio, H Henrik Ehrsson, and Alessandro Farnè. 2014. "Multisensory Representation of the Space Near the Hand: From Perception to Action and Interindividual Interactions." *The Neuroscientist* 20 (2): 122–35.

Bufacchi, Rory J, and Gian Domenico Iannetti. 2018. "An Action Field Theory of Peripersonal Space." *Trends in Cognitive Sciences*. 22(12): 1076-1090.

———. 2019. "The Value of Actions, in Time and Space." *Trends in Cognitive Sciences*. 23(4): 270-271

Butterfill, Stephen A, and Corrado Sinigaglia. 2014. "Intention and Motor Representation in Purposive Action." *Philosophy and Phenomenological Research* 88 (1): 119–45.

Chemero, A. 2009. *Radical Embodied Cognitive Science*. Cambridge: The MIT Press.

Cook, Erik P, and John HR Maunsell. 2002. "Attentional Modulation of Behavioral Performance and Neuronal Responses in Middle Temporal and Ventral Intraparietal Areas of Macaque Monkey." *Journal of Neuroscience* 22 (5): 1994–2004.

De Vignemont, F, and GD Iannetti. 2015. "How Many Peripersonal Spaces?" *Neuropsychologia* 70: 327–34.

DiCarlo, James J, and David D Cox. 2007. "Untangling Invariant Object Recognition." *Trends in Cognitive Sciences* 11 (8): 333–41.

DiCarlo, James J, Davide Zoccolan, and Nicole C Rust. 2012. "How Does the Brain Solve Visual Object Recognition?" *Neuron* 73 (3): 415–34.

---

[7] Thanks to Dan Burnston, Peter Clutton, Carl Craver, Frederique de Vignemont, Stephen Gadsby, Thor Grünbaum, Julia Haas, Ben Henke, Gabbrielle Johnson, Tony Licata, Ross Pain, two anonymous referees, and audiences at the Australian National University, Institut Jean Nicod, and Washington University in St. Louis for helpful comments and feedback. Work on this paper was supported by Australian Research Council grant FT140100422.

Funkhouser, Eric. 2006. "The Determinable-Determinate Relation." *Nous* 40 (3): 548–69.

Goddard, Erin, Colin Klein, Samuel G Solomon, Hinze Hogendoorn, and Thomas A Carlson. 2018. "Interpreting the Dimensions of Neural Feature Representations Revealed by Dimensionality Reduction." *Neuroimage* 180: 41-67.

Graziano, Michael SA, 2006. "The Organization of Behavioral Repertoire in Motor Cortex." *Annu. Rev. Neurosci.* 29: 105–34.

Graziano, Michael SA, 2016. "Ethological Action Maps: A Paradigm Shift for the Motor Cortex." *Trends in Cognitive Sciences* 20 (2): 121–32.

Graziano, Michael SA, G.S. Yap, C.G. Gross, and others. 1994. "Coding of Visual Space by Premotor Neurons." *Science* 1054–4.

Graziano, Michael SA, and Dylan F Cooke. 2006. "Parieto-Frontal Interactions, Personal Space, and Defensive Behavior." *Neuropsychologia* 44 (6): 845–59.

Graziano, Michael SA, Xin Tian Hu, and Charles G Gross. 1997. "Coding the Locations of Objects in the Dark." *Science* 277 (5323): 239–41.

Graziano, Michael SA, Lina AJ Reiss, and Charles G Gross. 1999. "A Neuronal Representation of the Location of Nearby Sounds." *Nature* 397 (6718): 428.

Graziano, Michael SA, Charlotte SR Taylor, and Tirin Moore. 2002. "Complex Movements Evoked by Microstimulation of Precentral Cortex." *Neuron* 34 (5): 841–51.

Graziano, Michael SA, Charlotte SR Taylor, Tirin Moore, and Dylan F Cooke. 2002. "The Cortical Control of Movement Revisited." *Neuron* 36 (3): 349–62.

Grush, Rick. 2007. "Skill Theory V2.0: Dispositions, Emulation, and Spatial Perception." *Synthese* 159: 389–416.

———. 2009. "Space, Time, and Objects." In *The Oxford Handbook of Philosophy and Neuroscience*. Oxford University Press.

Janssen, Annelli, Colin Klein, and Marc Slors. 2017. "What Is a Cognitive Ontology, Anyway?" *Philosophical Exploration* 20 (2): 123–28.

Klatzky, Roberta L. 1998. "Allocentric and Egocentric Spatial Representations: Definitions, Distinctions, and Interconnections." In *Spatial Cognition*, 1–17. Springer.

Lycett, Stephen J, Noreen von Cramon-Taubadel, and Metin I Eren. 2016. "Levallois: Potential Implications for Learning and Cultural Transmission Capacities." *Lithic Technology* 41 (1). Taylor & Francis: 19–38.

McCloskey, Michael. 2009. *Visual Reflections: A Perceptual Deficit and Its Implications.* Oxford University Press.

Noel, Jean-Paul, and Andrea Serino. 2019. "High Action Values Occur Near Our Body." *Trends in Cognitive Sciences* 23(4): 269-70.

Obrien, Gerard, and Jon Opie. 2009. "The Role of Representation in Computation." *Cognitive Processing* 10 (1): 53–62.

Pouget, Alexandre, and Terrence J Sejnowski. 1997. "Spatial Transformations in the Parietal Cortex Using Basis Functions." *Journal of Cognitive Neuroscience* 9 (2): 222–37.

Povinellia, Daniel J., James E. Reaux, and Scott H. Frey. 2010. "Chimpanzees' Context-Dependent Tool Use Provides Evidence for Separable Representations of Hand and Tool Even During Active Use Within Peripersonal Space." *Neuropsychologia* 48: 243–47.

Rescorla, M. 2009. "Cognitive Maps and the Language of Thought." *The British Journal for the Philosophy of Science* 60 (2): 377–407.

Ritchie, J Brendan. 2015. "Representational Content and the Science of Vision." PhD thesis, University of Maryland.

Rossetti, Angela, Daniele Romano, Nadia Bolognini, and Angelo Maravita. 2015. "Dynamic Expansion of Alert Responses to Incoming Painful Stimuli Following Tool Use." *Neuropsychologia* 70: 486–494.

Shea, Nicholas. 2018. *Representation in Cognitive Science*. Oxford: Oxford University Press.

Sterelny, Kim. 2003. *Thought in a Hostile World: The Evolution of Human Cognition*. Malden, Massachusetts: Blackwell Publishers.

Webb, Barbara. 2006. "Transformation, Encoding and Representation." *Current Biology* 16 (6): R184–R185.

Webb, Barbara, and Antoine Wystrach. 2016. "Neural Mechanisms of Insect Navigation." *Current Opinion in Insect Science* 15: 27–39.

Whittaker, John C. 1994. *Flintknapping: Making and Understanding Stone Tools*. University of Texas Press.

Wirth, Niklaus. 1976. *Algorithms + Data Structures = Programs*. Prentice Hall.