

MARK ZANGARI AND DAN CENSOR

SPECTRAL REPRESENTATIONS:

An Alternative to the Spatiotemporal World View

ABSTRACT. Is it possible to construct an alternative framework for the description of physical reality that is not based on space and time? According to Kant, because of the incorrigibility of the spatiotemporal scheme, the contents of any such alternative will be beyond our cognitive grasp. Nonetheless, the possibility of constructing such a descriptive scheme poses itself as an intriguing challenge. In this paper, we attempt to answer this challenge by exploiting an analytical tool extensively used by physicists and engineers: the utilisation of multiple representations of the same data using integral transforms – for example, the *Fourier transform*. According to the “spectral” picture we propose, frequencies and wavelengths, related to time and space by the Fourier transform, form an adequate alternative basis for describing physical systems, and one that is of equal richness to the spatiotemporal picture. Furthermore, the routine use of Fourier analysis applied to familiar physical systems suggests ways in which sense organs might be constructed such that they naturally give rise to a spectral representation of the world. Although we might never be able to imagine perceiving the world in this way, we can nonetheless describe the contents of such a mode of perception and relate them precisely to our own world-view. Hence, we suggest in a Kantian spirit, that the spatiotemporal description of the world is not unique.

1. INTRODUCTION

It has long been contended that our knowledge regarding the nature of space intimately depends on the epistemic status of geometry and on the ontic status of inferences from geometric propositions. Since the ancient Greeks, the extraordinary deductive richness of Euclid’s postulates and their almost self-evident character led to the belief that spatial properties necessarily conformed to Euclidean geometry and thus, were determinable *a priori*. It appeared that Euclidean geometry, due to its being a product of “pure” reason, provided a unique description of spatial properties that lay beyond the possibility of empirical falsification. The properties of space, it seemed, were uniquely fixed by logic alone and hence, the need for experimental investigation was obviated by deductive reasoning.

In the late eighteenth century, during an age when such ideas were coming under increasing attack from a growing empiricism, an important and subtle defence of the *a prioricity* of geometry and its consequences

Synthese **112**: 97–123, 1997.

© 1997 Kluwer Academic Publishers. Printed in the Netherlands.

for our understanding of spatial properties was constructed by Kant. In the *Critique of Pure Reason*, Kant identifies space and time, not as objectively real, but as transcendental ideals:

Space is a necessary a priori representation, which underlies all outer intuitions. We can never represent to ourselves the absence of space, though we can quite well think of it as empty of objects. It must therefore be regarded as the condition of the possibility of appearances, and not a determination dependent upon them. It is an a priori representation, which necessarily underlies outer appearances. [B 39, *The Critique of Pure Reason*.]

Thus, according to Kant, our ability to deduce facts about space and time *a priori* is understood not as a corollary of any irreducibly *ontological* feature of the world, but as an epistemic constraint rooted in the way we perceive. From this, it follows that, were our perceptive or cognitive apparatus different in some way, our physical description of the world might not be based on a manifold of spacetime events, but on an underlying structure that is *not* necessarily spatiotemporal. Adding flesh to such an assertion, however, immediately confronts an obvious practical obstacle, namely the inscrutability of this alternative underlying structure, whatever it may be. For, given Kant's claim, we cannot even conceive of the world in terms of anything other than spatiotemporal predicates. Note that although this restriction implies we will never be privy to the *content* of non-spatiotemporally ordered perception, we can nonetheless imagine the *possibility* of such a conception. Be that as it may, the claim lays down a tantalising challenge: is it possible to construct a physical representation of the world that does not rest on irreducible spacetime properties and relations?

Strawson in his 1959 monograph *Individuals* took up this challenge in his attempt to answer the question "Could there be a scheme, providing for a system of objective particulars, which was wholly non-spatial?" (p. 62). His approach was based on investigating to what degree we could develop a perceptual scheme in which objects external to ourselves could be individuated and identified after all senses except hearing are deleted. The deletion in question here is more severe than just sensory deprivation as not only is sensory input prohibited, but the very cognitive framework for accepting such input, were the senses functional, is absent. In Strawson's scheme, the richness of the spatial representation is no longer available to us and the only prospective replacement is the one-dimensional world of sound. Not surprisingly, the world-view that emerges from such a soundscape is considerably more limited than our own and suggests that spatiotemporal predicates probably are necessary if we are to retain a picture of an external world populated by identifiable, objective entities other than ourselves.

In this paper, we shall revive the quest to fulfill the Kantian challenge. However, our attempt to provide an alternative basis for ordering our description of the physical world differs somewhat from Strawson's, both in spirit and content. Rather than simply eliminating spatial predicates and attempting to use one of our remaining senses to identify objective particulars, we shall replace spatiotemporal concepts with a new conceptual framework, but one of similar richness to space and time. Our motivation comes from a suggestion implicitly raised by a manoeuvre that every physicist and engineer has performed, and yet with philosophical implications that hitherto have not been investigated to any great depth. This move is best exemplified by the physics of something like a vibrating string: the kinematics of the string can be mathematically described in terms of (at least) two distinct functions, one that conveys spacetime dependent behaviour, say $f(x)$, and another – its *Fourier transform* $h(k)$ – that represents the string's motion in terms of a frequency-wavelength spectrum.¹ Both convey the same information, although it is arranged very differently in each case, and there is no principled way of choosing either as more “correct”. This strategy, furthermore, is not limited to periodic or quasi-periodic systems like vibrating strings (although it is often most useful in such circumstances) but can be applied to (almost) *any* kinematical function in spacetime. As such, it appears to offer irreducible frequencies and wavelengths as a viable candidate for an “alternative”, non-spatiotemporal structure to underlie physics. We shall develop the details of this in Section 2 below and, assuming the strategy is feasible, speculate in Section 4 on why we have opted for space and time as our foundational notions in preference to other possibilities. Finally, in Section 5 we examine some possible arguments that we, as spacetime oriented creatures, might prefer in defence of our naive view that, despite the equivalence of the two descriptions, spacetime is nonetheless *ontologically* more fundamental.

Our analysis also sheds some light on why Strawson's sound world was unable to support the richness of perceptual distinctions that we can make in spacetime. As we discuss in Section 4, human perception of sound makes extensive use of the Fourier transform, but represents (mainly) *temporal* properties rather than spatial ones as might be suspected simply from dimensional considerations. Thus, sound is more properly associated with our sense of time, rather than our perception of space. In fairness to Strawson, he was not seriously proposing hearing alone as a substitute for the full spatiotemporal conceptual space currently at our disposal. Had he been, however, the attempt might be classified as a case of trying to fit the proverbial square peg into a round hole (or, more properly, a four-dimensional hypercube into a one-dimensional line). In the scheme we

propose, spatial properties are replaced by a wavelength spectrum of the kind we introduce in Section 2 below. In order to carry this out, we must define a way of translating a spatiotemporal description into the alternative we are proposing. Our strategy will utilise the transformation properties of the mathematical functions that represent spatiotemporally dependent properties. While such mathematical representations already constitute the language of physics, there is no reason why they cannot be employed in quantitatively describing the contents of our sensations as well. Consider vision, for example, which serves as our primary sense for individuating objects of perception and attributing spatial properties to them. The contents of our visual field can easily be represented mathematically by a function which specifies the intensity and colour of the light arriving at each of our eyes from any direction at any time. Obviously, hearing can be represented in a similar way. The power and accuracy of such a representation is amply demonstrated by mechanical devices, such as electronic cameras and microphones, which mimic the function of our sense organs and produce electrical output that is most naturally described by mathematical functions. What we seek, therefore is a transformation that maps the mathematical functions representing spatiotemporal descriptions to equivalent alternatives which no longer need be interpreted in terms of spacetime. Before proceeding, however, let us briefly outline some of the properties of the kind of transformation we need.

While much of the discussion of spacetime is couched in terms of transformation properties, these are generally restricted to those that map one representation of spacetime to another either by a change of coordinate chart (passive transformations) or by a modification of the intrinsic spacetime structures themselves (active transformations). Passive transformations are of particular interest to our cause because they allow us to change the representation of the objects being described, and this is just what we are setting out to do. But passive transformations do not have sufficient scope to change the representation to the extent that we require. For, in the way they are usually construed, passive transformations do not operate on the representation directly, but on the *mapping* between the intrinsic space and the abstract set that we use to describe it. We get different representations by using different mappings. In the particular case of spacetime theories, the intrinsic space is a four-dimensional topological manifold whose points are represented by correlating them to unique elements of \mathbb{R}^4 (via the coordinate mapping). To formally describe a spacetime dependent quantity, we typically use functions of the form $f : \mathbb{R}^4 \rightarrow \mathbb{R}$, whose arguments we physically interpret as spacetime points. And this is just where the standard idea of a passive transformation is too restrictive.

For, if we passively transform our representation, all we effectively do is change the correlation between points on the spacetime manifold and elements of \mathbb{R}^4 that represent them. Having done this, the arguments of the typical function $f(x)$ are still physically interpreted as spacetime points, but the 4-tuples of real numbers now simply correspond to different points on the manifold. We have changed the coordinates, but not the nature of what those coordinates are describing. Their physical dimensions both before and after a passive transformation are spatiotemporal.

What we seek is the ability to operate on the arguments of such functions directly in a way that may allow their physical interpretation to be altered. To coin a phrase, let us call such transformations “operative”, while we shall refer to the familiar passive transformations as “inoperative”, for the reason that they do not operate on the argument variables directly. Such changes of variable are of immense utility in solving differential and integral equations and, as we shall see in the following section, open the door to the possibility of translating a spacetime-based representation of the physical world to a non-spatiotemporal one. All we need is an operative transformation such that the new variables can be physically interpreted without the necessity of first reducing them to (relations among) spacetime points. Furthermore, if this mapping is non-singular (or invertible), then we can be sure that anything which can be represented by a given function can also be represented by its image under the transformation, and vice-versa. Thus, invertibility ensures that the translation between the two schemes is “two-way” for any function. If such a transformation is achievable, it might be possible that we have a plausible alternative to the spatiotemporal mode of describing the physical world. But *is* there an operative transformation that fulfils this tall order? That is, one that operates on *functions* of \mathbb{R}^4 , is invertible, and replaces the variables of the functions in its domain with new quantities whose physical interpretation need not ultimately rest on spatiotemporal primitives.

2. THE FOURIER TRANSFORM

Fortunately, our search need not carry us to places as strange and exotic as we might have initially expected. In fact, the analysis of simple harmonic oscillators, such as vibrating strings provides a paradigm example of the kind of transformation we are after. For, suppose we have a system executing some sort of periodic motion, say a violin string shortly after being plucked. We can easily represent a spatial snapshot (limited to one dimension for simplicity) at any given time by a function $f(x)$ or alternatively, the temporal behaviour of any point on the string by a (different) function

$f(t)$. However, it is well known that we can also represent the physical state of a vibrating violin string in terms of a *frequency* (or *wavelength*) *spectrum*. That is, the motion of the string can be uniquely produced by the superposition (or linear combination) of sinusoidal oscillations at different frequencies (or wavelengths). These frequencies are often called the “fundamental” and “overtones” that give the violin (or any instrument, for that matter) its unique timbre. The unique distribution of frequencies corresponding to the particular motion of the string is usually represented by a function we call the *frequency spectrum*. This specifies the amplitude or “strength” of the sinusoidal oscillation at each frequency required to build up the overall motion of the string. Let us represent the frequency spectrum by the function $h(\omega)$ where ω , the *angular frequency*, is related to the frequency ν via $\omega = 2\pi\nu$ and therefore has the dimensions of inverse-time. Alternatively, we could specify the wavelength spectrum $h(k)$ where k , the wave-number, is inversely related to the wavelength λ via $k = 2\pi/\lambda$. Therefore, k has the dimensions of inverse-length and is a measure of a quantity that Censor (1991) refers to as “nearness”.

Summing the contributions of all the oscillations weighted by $h(k)$ over all wavelengths produces the required spatial behaviour of the string (similarly, weighting by $h(\omega)$ and summing over all frequencies reproduces $f(t)$, the time-dependent description of the motion). Of course, this will work for any periodic system (and, it will turn out, for any physical system that can be described in spacetime by reasonably well behaved functions), as is readily illustrated in another equally ubiquitous example: the production of white light. Although most people would be hard pressed to describe the spatiotemporal wave pattern of white light, any high-school student could (hopefully) tell you that the colour white results from the superposition of light of all frequencies,² as can easily be demonstrated by passing it through a glass prism.

This provides some motivation for, and gives us an informal account of what is, in effect, an operative transformation. In going from the spacetime to the spectral representation, we are defining a transformation from the function $f(x)$ (or $f(t)$) to $h(k)$ (or $h(\omega)$). However, we haven’t given the formal rule for the transformation yet, nor have we determined whether it is one-to-one and hence, invertible. Let us turn to this now. Monochromatic oscillations of wave-number k (or angular frequency ω) are usually represented by complex exponentials of the form e^{ikx} (or $e^{i\omega t}$), since the real and imaginary parts of this expression are $\cos kx$ and $\sin kx$ ($\cos \omega t$ and $\sin \omega t$), respectively. For consistency with what follows, for the moment we will choose to represent our monochromatic vibrations by e^{-ikx} , but the reversed sign in the exponent changes nothing of any substance. (Note

that whatever applies for k and x also applies for w and t . To make what follows less repetitious, we shall not restate this in each case, but assume it unless otherwise stated.) We have already stipulated that $h(k)$ is the amplitude of the contribution from the simple oscillation of wave-number k required to produce the motion described in $f(x)$. To reconstruct $f(x)$, we just have to “sum” over $h(k)e^{-ikx}$ for every k . But k is a continuous variable so instead of summing, we must integrate. That is, we can write the position-dependent function $f(x)$ in terms of the spectral function $h(k)$ according to:

$$(1a) \quad f(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} h(k) e^{-ikx} dk$$

Equation (1a) is the well known Fourier transform³ of a function of one real variable and possesses the associated inverse transformation:

$$(1b) \quad h(k) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} f(x) e^{ikx} dx$$

Note that the factors of $1/\sqrt{2\pi}$ are introduced to normalise the integrals in (1a) and (1b) symmetrically. Note also that, because the Fourier transform is invertible, the mapping between the domain containing $f(x)$ and the co-domain containing $h(k)$ is one-to-one. Thus, every function $f(x)$ has a unique Fourier transform and vice-versa. It is common in mathematics texts to refer to the first equation (1a) as the inverse transformation and second one (1b) as the “forward” one. By examining both parts of Equation (1) however, it should be clear that the choice of either (1a) or (1b) as the “forward” transformation (and the other one as the inverse) amounts to no more than specifying the sign in the exponent. But, because of the high degree of symmetry between the two equations, this is largely a matter of convention and either definition suffices.

Both parts of Equation (1) can readily be generalised to higher dimensions, in particular to encompass four-dimensional spacetime functions. Instead of $f(x)$, we could easily consider $f(t, x, y, z)$ and correspondingly, instead of $h(k)$, we would then employ $h(\omega, k_x, k_y, k_z)$, where k_x, k_y , and k_z are, respectively, the x, y , and z components of the wave-vector (the higher dimensional generalisation of the wave-number). Since the four variables in both f and h are independent, we could define the four-dimensional Fourier transform simply by multiple integration. The exponent in the integrand is modified so that the desired function (f or h) is multiplied by $e^{\pm i(\omega t + k_x x + k_y y + k_z z)}$ and the normalisation constant of $1/\sqrt{2\pi}$ is raised to the power of 4. However, because of the above

mentioned symmetry in the forward and inverse transformations, and the resulting conventionality in the choice of sign in the exponent, we are really free to choose the sign of *each term* in the exponential *independently*. Thus, we could if we so wished (and we do, for reasons that will be apparent immediately) define the exponential in the four dimensional Fourier transform to be $e^{i(\omega t - k_x x - k_y y - k_z z)}$. A consequence of this form is that the real and imaginary parts of the exponentials (i.e. $\cos/\sin(\omega t - \mathbf{k} \cdot \mathbf{x})$) are simply the equations for monochromatic plane waves filling spacetime. Thus, the reasoning from our one-dimensional example of the violin string (or white light) carries through to the four-dimensional case, with the one-dimensional sinusoidal oscillations in the former replaced by monochromatic plane waves filling spacetime in the latter.

Before proceeding, we must draw attention to one further important observation. We initially introduced the Fourier transform by attempting to describe *periodic* systems in terms of sums of sinusoidal oscillations of definite wavelength. And, one might feel that this is reasonable, since the periodicity of a vibrating string, or oscillating electromagnetic field lends itself to this kind of description quite naturally. However, there is no mention made of the periodicity of f or h in either of Equations (1a) or (1b). Nor, indeed, should there be. The Fourier transform performs equally well operating on highly periodic functions, or on highly non-periodic ones. A moment's reflection should satisfy the reader that this is the case. For the closer the function f is to being "perfectly" periodic (i.e. a pure, monochromatic sinusoid), the fewer wave modes are needed to describe it and the more "peaked" h will be around the value of k corresponding to the wave-number of f . This means h itself is highly non-periodic. But if h exists, so does its Fourier transform (since h is the Fourier transform of f and this must be invertible). Because of the symmetry in Equations (1a) and (1b), we could then swap the roles of f and h (or swap our sign convention in the exponentials) and the same relation between them would hold – f , the spacetime function, would be "peaked" and h , its spectral equivalent would be close to a sinusoid. Thus, arbitrary non-periodic functions can be represented spectrally in terms of plane waves, so long as the integrals in Equation (1) exist. We should note that we are not implying that the Fourier transform possesses any special formal properties over and above any other invertible, linear integral transform. However, for our present purposes, it comes with an added bonus: the primitive entities which replace spacetime points in the transformed scheme (i.e. wavelengths and frequencies) are easily physically interpretable in their own right (albeit usually in terms of spacetime points, but we shall return to this point later). Hence our choice.

3. SPECTRAL VS. SPACETIME

Our somewhat mathematical interlude finally over, let us now examine what we have to show for our labours. The Fourier transform provides a means of translating either way between a formal physical description based on functions of spacetime, and an alternate one based on spectral functions. In the former, physical quantities are expressed as relations between magnitudes (in this case, field strengths, trajectories – really a special case of a field – etc.) and spacetime points. In the latter, physical quantities are expressed as relations between magnitudes (this time, spectral distributions) and plane wave modes. Note that this highlights one of the key properties of operative transformations mentioned above, namely that following an operative transformation, the physical interpretation of functions and their arguments may undergo a radical, qualitative change, even though the same “physical reality” is being described. Indeed, it was this very feature that we hoped might allow us to construct a framework for describing the physical world that did not ultimately rest on spatiotemporal predicates. What we are suggesting here is that the Fourier transform provides a possible gateway into such a description.

One might immediately balk at this idea with the objection that frequencies and wavelengths are *not* irreducible and *do* ultimately rest on the spatiotemporal framework. After all, frequency is expressed in terms of “cycles per *second*” and wavelength as the “*distance* between equiphase *points*”. However, such a naive response only begs the question. The transformation between the spacetime and spectral descriptions in Equations (1a) and (1b) are invertible and almost identical in form. Any (suitably integrable) function that can be written down in one domain has a unique, well-defined counterpart in the other. Formally, there is nothing in Equation (1) to suggest that $f(x)$ (say) is somehow “more fundamental” than $h(k)$, or that the former is the “true” representation of the physical system being described while the latter is some kind of artificial construction. If we want to run this kind of argument, then appealing to the units we use in measuring the physical quantities amounts to no more of a justification than simply asserting the conclusion. For we employ spacetime dependent units for frequency and wavelength precisely *because* we take space and time to be fundamental, not the other way round, so the kind of argument suggested above amounts to no more than a stipulation. Because of the symmetry between (1a) and (1b) and their invertibility, we could just as easily reverse the dependency relation between k and x . That is, we could take time and space to be derived from (irreducible) frequency and wavelength. So, instead of constructing standard rods and clocks on which our measurements are based, we could define a plane wave with

a standard phase. In that case, time intervals would be derived quantities expressed as, say, “cycles per standard oscillation”, as too would spatial intervals, these being expressed as, “cycles per standard wave-number” or something similar. Hence the type of support mooted above for privileging spacetime descriptions over spectral ones fails.

So where does our deep-seated belief that a spatiotemporal description is somehow more fundamental than a spectral one come from? We will put forward a case in Section 4 that this preference may be an accidental feature of the way our sensory perceptual mechanisms evolved. Indeed, we can easily posit the existence of creatures whose sense organs are constructed differently from ours with the consequence that their subjective judgement about which description is more fundamental is reversed. That is, the “objects” they perceive in the world are individuated by their spectral properties. One may feel that this is just a way of speaking, as a collection of similar spectral properties, however identifiable and enduring, is not deserving of the epithet “object”, or at least far less so than something with closely knit spatiotemporal properties. However, ruling out the viability of the spectral world view simply by insisting that objects be spatiotemporally defined is again either a vacuous stipulation or another case of question begging. Our meaning of “object” in the present context is therefore close to Strawson’s “particulars distinguished by the thinker, etc.” (1959, 61). For the spectrally-oriented creatures (with a level of intelligence comparable to ours) that we shall meet in Section 4, “spatiotemporal” talk would be just a convenient if rather abstract way to represent (spectral) objects that display certain features (namely, those whose spectral properties have regular, periodic patterns). But this is looking ahead. For the present, let us assume that treating frequencies and wavelengths as the foundations of a conceptual scheme is feasible, and concern ourselves with further details of the spectral representation.

As we have repeatedly been stressing, the Fourier transform is invertible and hence defines a one-to-one mapping between the spacetime and spectral representations. Thus, any proposition containing a given spacetime predicate has a unique counterpart containing a spectral predicate (and, of course, vice-versa) and, if the transformations are carried out in a consistent manner, the truth values of such propositions should be preserved. In this sense, the amount of information contained in either representation is the same. Given that this is the case, one might wonder why anyone would bother performing a transformation of this kind, even though the Fourier transform is one of the most powerful analytical tools available to physicists and engineers. A possible answer is economy, in the sense we shall now describe. Although all real physical systems are larger than

idealised points, some systems behave very nearly as if they were points and representing them this way proves to be very efficient. Alternatively, other systems (fields, in particular) are not very point-like at all and must be represented as some kind of spatiotemporally extended object. An idealised limiting case of this is the monochromatic plane wave: this extends throughout all space eternally. Of course, this is every bit as much of an idealisation as the dimensionless point and real physical systems fall somewhere between these two extremes. If a given property of an object is spatially localised, then a spatiotemporal description of it is very economical; we need only specify the components of the four-vector x . If, on the other hand, a property is spatially distributed and, in particular is wave-like, we must specify infinitely many x values. In the spectral representation, the situation is reversed. Point-like objects require infinitely many superimposed wave modes to describe them, while wave-like objects require only the specification of the four components of the wave four-vector k . Thus, depending on the characteristics of the system we wish to describe, one or other mode of representation may prove to be more economical and will naturally highlight certain features. Of course, in between the two extremes, there are objects (such as non-periodic, non-localised fields) which are about equally complicated in both regimes. Typical of such systems are those described by Gaussian distribution functions; the Fourier transform of the Gaussian $\exp(-\frac{1}{2}x^2)$ is just itself (with x replaced by k and modulo normalisation constants).

The upshot of all this is that there is, formally at least, an alternative to the “space” (used here in a mathematical sense to mean something like “set” or “domain”) in which physical descriptions are founded – it can be consistently formulated as “ k -space” (or “spectral space”) in addition to the more familiar “ x -space” (or “spacetime”). We shall reserve our comments about the ontological questions this raises to the final section. For the present, let us discuss some further features of k -space. One of the key defining features of a spacetime is the way that physical descriptions transform from one observer to another. In our best spacetime theory, these transformations are (locally) Lorentz transformations and we regard this as a fundamental, defining feature of the kind of spacetime we are dealing with. But if our claim that we can consistently represent our physics in k -space is to hold, we must be able to translate the Lorentz transformation into a unique spectral counterpart. For, since spacetime is a derived notion in k -space, spatiotemporal mappings are not transformations of primitive entities, but are a consequence of the transformation of fundamental *spectral* properties, i.e. mappings between frequencies and wavelengths. So perhaps here is the stumbling-block of our argument; perhaps there is no

sensible way to translate the Lorentz transformation to something fundamental in k -space. Or, put in the form of a question, can we “derive” the Lorentz transformation in spacetime from something fundamental in spectral space?

The answer is an emphatic “yes”. In a spectral picture, Lorentz invariance is represented by the relativistic Doppler effect (Doppler 1842; Gill 1965). Comparing this and the standard, spacetime Lorentz transformation by velocity⁴ v (in units where $c = 1$, which will be assumed from here on) in one dimension shows the close formal similarity between the two:

$$(2) \quad \begin{array}{ll} \omega' = \gamma(\omega - vk_x) & t' = \gamma(t - vx) \\ k'_x = \gamma(k_x - v\omega) & x' = \gamma(x - vt) \\ k'_y = k_y & y' = y \\ k'_z = k_z & z' = z \end{array}$$

where $\gamma = 1/\sqrt{1 - v^2}$. Normally, the Doppler effect is derived as a consequence of the Lorentz invariance of relativistic spacetime. However, in the k -space representation, the situation is reversed; the Doppler effect is primitive and from it, the Lorentz spacetime transformation can be derived (by running the usual derivation in reverse). Which of the two we take as fundamental depends on whether our description is grounded in x - or k -space. And, in a fashion comparable to both parts of Equation (1), the high degree of similarity between the Lorentz and the Doppler transformations singles out neither as obviously dependent on (or a consequence of) the other (further to this, see Censor 1991a and 1991b).

In this regard, it is interesting to note that the relativistically correct results for reflection by a moving mirror were actually derived by Abraham (1904), *before* Einstein’s (1905) paper introducing special relativity, and quite independently of it. Unfortunately the whole impetus behind Abraham’s work concerned the solution of the specific problem at hand and, not possessing the retrospective insight we can apply, Abraham did not recognize it as heralding an overhaul of the most fundamental concepts in physics. Perhaps it is *because* his intuitions were so firmly grounded in spacetime (as are ours) that Abraham did not appreciate the revolutionary and highly unexpected consequences of his discovery – these not being recognised until the theory was later explicated by Einstein in terms of space and time (see also Pauli 1958, who mentions many of the relevant studies). So, in the context of the present discussion, not only is it the case that special relativity *could* be couched in terms of frequency/wavelengths with the Doppler effect but that, thanks to Abraham, in one sense it actually *began* that way!

Although the ready intertranslatability of the spatiotemporal and spectral world views implies the kind of symmetry between the two descriptions that we have been exploiting, there is unquestionably a compelling *asymmetry* in the way we regard them. The spatiotemporal description is generally taken to be fundamental and we appear to have almost no choice but to represent the physical world to ourselves in this way; Kant's point was just that any other kind of representation is literally unthinkable. Although this still allows us an alternative mode of *description*, i.e. k -space, we can only make physical sense of such a picture by regarding it as a translation from the "real" spatiotemporal one. That is, by regarding it as merely an abstract representation of physical processes that we incorrigibly believe are *really* taking place in spacetime. But as we have seen, this asymmetry has no basis in the translation between the two regimes as there is nothing to single out either direction of the transformation. Perhaps, then, one or other representation has some intrinsic property that would enable us to settle the question of which is more fundamental (presumably in favour of x -space). But, other than our innate proclivity to interpret spacetime predicates "literally", which is what we are trying to justify or at least, to explain, there seems to be no such property. It is as if we are trying to assert that a description of some physical process written down in English represents something more fundamental than one written down in Hebrew on the grounds that we can only understand the Hebrew text by first translating it into English via well defined rules. Yet we are only in this position because we are native (and monolingual) speakers of English unable to extract any linguistic content directly from the Hebrew. Similarly, it seems that our perception allows us an acquaintance with space and time that is not extended to such things as frequencies and wavelengths, and it is from this that our confidence in the ontological primacy of space and time springs.

4. THE ROLE OF PERCEPTION

Our continual appeal to the equivalence of the x - and k -spaces does not diminish the fact that in most cases, our subjective experience of spatiotemporal properties seems far more immediate and irreducible than is the case for spectral ones. The question then becomes: is this subjective preference of the spatiotemporal over the spectral sufficient to imply the uniqueness of the former as a mode of perception, or even as an objective feature of the world? Or could it be that accidental physical constraints imposed by our sensory system lead us to favour representations based on

space and time and, were our sense organs and brains slightly different, might we find spectral representations more “natural”?

In order to answer this, let us begin by considering the role of our sensory apparatus in the perception of space. First of all, our perception of distance, unlike time, is not “direct”. To perceive time (however unreliably) seems to require merely that we are conscious, and no further input from the external senses appears to be needed. However, our primary perception of spatial properties is by virtue of our sense of vision, or our eyes’ capability to form images on the retina by focusing the light incident on the cornea.⁵ Because the pupil is very large compared to the wavelength of the light our eyes are sensitive to, the image formation process can be considered entirely within the realm of geometrical optics. This means that the wave nature of light which leads to diffraction and interference phenomena (or “physical optics”) can be ignored. The ray approximation for the propagation of light represents only its space-time trajectory and hence, the imaging mechanism of our eyes (setting aside colour perception, for the present) responds only to the positional information carried by the light.

However, if the size of the pupil were of the order of the wavelength of the light to which the retina is sensitive, then the imaging process in the eye would not lie in the realm of geometrical optics, but would be based on diffraction. In the case where such an eye operated using far-field (or Fraunhofer) diffraction, then the patterns produced on the retina are essentially the Fourier transforms of our geometric-optics images. Therefore, a creature endowed with such an eye would have its visual sensations ordered not by their spatial properties, but by their wavelengths.⁶ Such a mechanism provides the physical basis for a sensory system that generates a representation related to our spatial one by the Fourier transform. As an aside, let us note that our own senses already use the Fourier transform extensively. Our ability to perceive the pitch of a sound is due to the structure of our inner ear. The cochlea, which is a fluid filled snail-like coiled organ, acts much like a musical wind instrument in reverse, resonating differently at different sound pitches. In a woodwind instrument, pitch is controlled by the player covering or uncovering precisely positioned holes in the body of the instrument, thus changing the position of high pressure nodes in the vibrating air column and consequently, the frequency of the sound produced. In the inner ear, different frequencies produce pressure nodes at different locations. Along the cochlea there is a partition called the Basilar membrane, which is covered with hair cells. The pressure acting on the hair cells causes the attached nerve endings to fire. Thus the location of the pressure points on the Basilar membrane corresponds to the sound pitch perceived by the brain.

Consequently, our ears resolve a complicated time-dependent motion (the pressure fluctuations in the air) into its component frequency modes and we hear sound not as time-fluctuating air pressure, but as a superposition of frequencies.⁷ Yet, because of our attachment to the spatiotemporal world-view, our physical account of the mechanics of sound is still in terms of time-varying air pressure fluctuations (we shall return to this in more detail shortly). In a similar way, our ability to perceive colour (i.e. light frequency) arises from the three types of cone cells in the retina being (somewhat loosely) tuned to the frequencies of red, green and blue light. This time, there are only three modes, but they are far from “pure” as there is considerable width and overlap in the excitation spectra of the three types of receptors.

Returning to our main theme, we might wonder how the world would look to a creature endowed with a physical optics eye. Perhaps the best way to imagine this is to draw an analogy between our eyes vs. the creature’s, and a photographic camera vs. a holographic one. The workings of our eyes are not unlike that of the photographic camera, with the image on our retina not unlike that exposed on the camera’s film; similarly, the image on the “retina” of the creature’s eye is not dissimilar to the pattern on a holographic film. It is well known that such patterns contain all the information present in the spatial image, to which anyone who has viewed the reconstruction of three-dimensional “holograms” from such patterns will be able to attest. But our creature does not “reconstruct” the spatial image, nor has it any need to, given that all the information in one representation is also present in the other (recall that the translation between them is invertible). Because its ocular senses evolved using diffraction (perhaps because it inhabits a highly amorphous environment where an efficient representation of spectral properties is more useful than spatial ones), it seems reasonable to suggest that this creature’s perception of the world is based directly on this representation. Using the “holographic” image, it is perfectly able to navigate its way through any environment we could negotiate; indeed, to such a creature, the world *is* the content of this image, just as it is the content of our visual sensations to us.⁸ Although we and this creature may differ in our identification of objective particulars (to return to the language of Strawson 1959) in that what is a highly localised “object” to one will be a highly dispersed and possibly unconnected set of sensations to the other, the two representations are nonetheless commensurable and have identical information content. They are simply optimised to the efficient representation of different kinds of environments. Any question about the physical world that we can answer

can be translated into a question that the creature could answer and our answers would be the same.

The analogy between humans and the spectral-eyed creature becomes somewhat more difficult to explicate for the case of time sensation because the subjective experience of temporal becoming is not easily located in any given organ of our bodies. Consequently, we will need to do more work as it is not immediately obvious how we should tinker with familiar sense organs so that our creature, endowed with suitably modified equipment, will perceive some kind of frequency spectrum where we perceive temporal intervals. But, as we noted earlier, we humans already possess temporal Fourier transformers in our ears so let us use this as a starting point. Before we begin modelling alien sense organs on our aural receptors however, we must highlight a crucial difference between the “partial” spectral representation offered by our ears, and the “complete” spectral representation we are attributing to our spectral friend. As we shall see, however, what we need to consider in making this distinction ultimately leads us to some idea of how we might understand the replacement of temporal intervals by a frequency spectrum.

To ground what follows in a concrete example, let us suppose we are listening to some music, say a recording of Beethoven’s Pastoral Symphony. To avail ourselves of a visual representation of the sound, let us suppose that our hi-fi is equipped with a spectrum analyser which visually displays the frequency distribution of the sound as it changes in real time. Suppose also that we plug an oscilloscope, capable of storing its input in memory, into the output channel of the amplifier. As the music plays, the oscilloscope displays and records the time-dependent waveform of the sound coming from the loudspeakers, while our ears convert this into a time-dependent superposition of frequencies. The spectrum analyser does essentially the same thing as our ears but displays the result visually. And here is the key: the spectral representation of the sound from both the spectrum analyser and our ears is *time-dependent*. We hear the music as a *sequence* of pitches, harmonies, etc., and the lights on the spectrum analyser form patterns that vary *in time*. So time has not been eliminated from this representation of the sound to be replaced by frequency, as the spectral picture would have it. Rather, what our ears and the spectrum analyser do is act only on frequencies *within a given range* – in both cases, audible frequencies between about 20 Hz and 20 kHz. Therefore, any frequency component with a periodicity longer than about one-twentieth of a second will not show up on the analyser’s display, but will vary from one in the sequence of displays to the next. Equivalently, we will not hear

it as a pitch, but as a sequence of sounds progressing one after the other in time.

To illustrate this more clearly, imagine you are hearing a rapid-fire sequence of short, sharp clicks, but not so rapid that you can't distinguish them. Now steadily increase the rate of the clicking. Soon, you can't distinguish the individual clicks and, as the frequency goes up, there comes a point at which you stop hearing them as clicks but hear the sound as a continuous tone. The fact that this occurs indicates that the clicks have just reached a high enough frequency to trigger the lowest frequency detectors in your ears. Anything of lower frequency will sound sequential, like the sequence of clicks. Note that this frequency also roughly corresponds to the smallest time interval we can perceive; we do not sense that time passes between each vibration of the tone, for it is oscillating too quickly for our time sensor (however that is physiologically implemented) to keep up with. Indeed, it would take an immense amount of brain power to sample sound waves up to 40,000 times per second and process them in real time, which is what we would have to do were we to make sense of what we hear in this way. Instead, our sensory system evolved to optimally exploit the economy of the spectral description that we mentioned above in Section 3, when it comes to representing approximately periodic systems. Sound waves are quite periodic on time scales of one-twentieth of a second or less, but not so on longer scales – say a minute, or a year. Thus, over short time scales, we use the spectral representation while for longer periods, our perception of sound, just like all else we perceive, is temporally sequenced.

So the kind of spectral representation constructed by both the spectrum analyser and our ears is only partial as it excludes contributions to the frequency spectrum below a particular cut-off frequency. As a rough guide, the order of magnitude of the time interval that we must integrate over in Equation (1b) to do the translation from x to k -space is inversely proportional to the lowest frequency in our spectrum. Thus, our ears and the spectrum analyser only take in about one-twentieth of a second of Beethoven's Pastoral Symphony at a time to give the spectrum of frequencies above 20 Hz.

Now we turn our attention back to the oscilloscope and imagine that we print out a graph displaying the contents of its memory after the completion of the first movement which, on this particular recording, happens to last for eight minutes and fifty-seven seconds. What we would then have is a record of the time-dependent wave-pattern for the whole first movement of the Symphony. We could then reconstruct the time-sequence of displays on the spectrum analyser by Fourier transforming the first one one-twentieth of a second's worth of the graph, displaying the part of the result that

falls between 20 Hz and 20 kHz, and then going on to do the Fourier transform of the next one-twentieth of a second, and so-on, 10,740 times (or thereabouts) till we get to the end of the graph. This, in effect, is roughly what our ears and the spectrum analyser are doing in real time as the music plays, thus producing the time-sequence of frequency spectra that we see and hear.

But suppose we now invest in better equipment and buy a spectrum analyser that is sensitive to frequencies down to 2 Hz. To represent the music spectrally, we will only need 1,074 samples (i.e. one-tenth as many as before) and each one will represent the spectrum of roughly a half-second slice of the music. That is, we have one-tenth as many slices with each one ten times longer than before. Then (in principle), we could keep going and get better and better machinery to detect 1 Hz, 0.1 Hz, 0.01 Hz, 0.001 Hz, etc. by which time we could Fourier transform the whole first movement at once. Suppose also that we had evolved a little differently so that the lower cut-off frequency of our ears was also 0.001 Hz and this corresponded to our perceptual “sampling rate”. Consequently, we would have to wait 1000 seconds (or 16 minutes and 40 seconds) to build up each spectral distribution that we hear. As before, let us assume that we do not detect the passage of time during intervals shorter than the period of the lowest frequency sound our ears can detect. The first movement of Beethoven’s Sixth Symphony would then sound like an almost instantaneous “click”, as its entire duration is shorter than our perceptual sampling rate. However, despite having the whole symphony compressed into a single perceptual event, it does not necessarily follow that we have lost any information that we had before. Consider that we can tell the difference between a short burst of sound from a violin and one from a bass guitar despite not representing the time-dependent details of the sounds in a time dependent way. Instead, we encode them spectrally as each of the instruments’ unique timbre. If we drop our lower cutoff frequency by several orders of magnitude, and increase our discriminatory abilities accordingly, we could imagine a much wider range of possible “timbres”, one of which corresponds to the first movement of Beethoven’s Sixth Symphony and another of which corresponds to Led Zeppelin’s *Stairway to Heaven*, etc., etc.

If we keep lowering the cutoff frequency, we would eventually get to the stage where we would have to wait our entire lifetime to gather enough data to construct the complete frequency spectrum right down to the lowest frequencies. Then, we would no longer perceive a time sequence of spectra, but just one single frequency spectrum encoded into which are the details of all the time-varying phenomena we have witnessed throughout our life. In that case, we would have abandoned the use of time order and our mode

of representation would be completely spectral. Since our other senses can also be characterised as a set of time-varying intensities, this spectral representation is not limited to hearing, but could be extended to all other forms of perception. This may be the closest we can get to understanding how our hypothetical creature might perceive a spectral distribution instead of temporal intervals.

The very alien, almost unfathomable nature of such a notion reinforces the central role that time plays in our cognition. Space may also be seen as indispensable, despite the relative ease with which we could compare our sensations of distance with the creature's. The earlier comparison was possible only because we could represent the Fourier transformed images spatially (i.e. as a holographic interference pattern). Without resorting to the familiar representation, the Fourier transformed images would be as unimaginable to us as a cognitive sense based on frequency rather than time. Yet our inability to comprehend, or even imagine the subjective experience of seeing the world this way, while limiting, seems no more problematic than our analogous ignorance of what it is like to derive perceptual content from any other sense we do not possess – echolocation (i.e., “sonar”) for example, or the ability to sense the electric fields emitted by the muscles of animals (for a lovely account of such things, see Dawkins 1991, chapter 2). Just as we can no more know what it is like to be the spectral-based creature than we can know what it is like “to be a bat” (Nagel 1979), we can no more deny the legitimacy of the spectral creature's perceptual content than we can for a bat.

The upshot of all this is that the form our perception takes seems not to be unique, but can easily be imagined to be structured very differently and to depend on different physical mechanisms, as our spectral creature's “physical optics eye” and “super-low frequency ear” demonstrated. It appears to be, in effect, an accident due to the peculiarities of our evolutionary history. If this is the case, however, a rather unsettling consequence seems to follow. For, suppose we got into an argument with a spectral-eyed creature about *ontology*. An alternative, equivalent representation of the world is something we can accommodate, but the lack of any arguments, other subjective, speciesist ones, to privilege one of the representations as “the real thing” hints at the sort of underdetermination that has often led to some kind of eliminativism. That is, it suggests that if there is nothing to objectively single out one among a number of representations, then whatever distinguishes those representations cannot be part of ontology and the distinction amounts to merely the stipulation of some convention. A paradigm example of this is the equivalence of all Galilean frames in Special Relativity, and the resulting elimination of the absolute rest frame from its

ontology. So where would our ontological argument with a spectral-eyed philosopher lead?

5. ONTOLOGICAL ARGUMENTS

Let us suppose for a moment that our argument led to a stalemate. That is, neither side had a good enough case to privilege either perceptual scheme as reflecting the underlying structure of the world. Since both schemes then amount to equivalent descriptions analogous to a change of Galilean reference frame in relativity, we would be encouraged by such people as Reichenbach (1927) and Poincaré (1905) to deny the reality of whatever it is that distinguishes them. But precisely *what* is it that distinguishes them? In the case of special relativity, the answer was fairly straightforward - we remove the notion of absolute velocity and replace it with an absolute notion of affineness (or “straightness”. This is how inertial world-lines are characterised). But *what* structure that we currently take to be objectively real would we have to remove if our arguments for the privileged status of x -space descriptions do not triumph over those of our spectral opponent (nor vice-versa), and what would we replace it with? This is a difficult question to answer, or rather, the answer is difficult to articulate. For, by usurping space and time, we are denying recourse to the most basic language we use to describe the world. To return to our earlier analogy, it is as if we have just discovered that neither English nor Hebrew is the “one true language” of which all others are merely translations. Hence we cannot provide an adequate semantics for Hebrew, say, by asserting that the model of a given Hebrew sentence is its English translation (recalling that in our case, translations are unique via Equation (1)). Rather, both English and Hebrew are modelled by a deeper structure which must be specified using some kind of metalanguage of which none of us are native speakers (but some of us might have learnt it formally). In our case, the mathematical “metalanguage” is functional analysis, since both the spectral creature and ourselves will be able to construct identical function spaces to describe the world. Furthermore, since the Fourier transform is an isomorphic mapping of a function space onto itself, each would be able to interpret the opposing world view as simply a transformation of the arbitrary basis for the underlying space. But this is leading us into mathematical complexity that we wish to avoid. Let us leave this point by stating that physicists are already used to formally describing the physical world in terms of function (or Hilbert) spaces – indeed, this is the central pillar of quantum theory. But neither they nor we can make any sense of this description without first translating it into

a spatiotemporal one, and thus creating for ourselves the interpretational dilemmas of quantum theory.

But before resigning ourselves to this fate, however, we should ask whether we need give in to the spectral philosopher in the first place. For despite his/her/its persuasiveness and eloquence, it is possible that we might have superior arguments at our disposal allowing us to proclaim the spacetime representation as objectively more fundamental than its spectral translation. To investigate this further, we must revisit precisely what is going on in the translation from the spatiotemporal to the spectral in order to highlight an aspect that might allow us to score a point in the argument against our alien philosopher. That point is that in going from x -space to k -space as in Equation (1b), the integration variables all range from $-\infty$ to ∞ (everything that follows also holds for translating in the opposite direction, i.e. as in Equation (1a)). That is, to produce the spectral representation, we must be able include the contributions from *all positions in space at all times*. Relativity aficionados will immediately protest that this requires access to information from outside our light-cone, something that relativity rules out. But this objection can easily be met by modifying our translation procedure a little. Specifically, we can cure the problem by multiplying the integrands in both parts of Equation (1) by the characteristic function of the light cone associated with a given point. That is, multiply by a point-function that equates to 0 if a point lies outside the given light cone and 1 otherwise. Thus, contributions from inside the light cone are not affected and those from outside are switched off.

But there is a second, potentially more unsettling problem that is not so easily fixed. In order to translate the time variable t to angular frequency ω , we must include contributions from all times, *past and future*. But it appears that however we humans perceive time, we have limited access to the past (via memory) and no access whatever to the future. This poses somewhat of a problem for constructing the kind of representation we are attributing to the spectral creature. There are two possible approaches one could take to this dilemma and which seems more attractive depends on one's ontological bents regarding the nature of time. The first route would be favoured by those sympathetic to so-called "four-dimensionalism", i.e. the idea that the universe is an integrated four-dimensional object and our apparent sense of the "passage" of time is merely a subjective illusion that will probably be explained by a better understanding of neurophysiology (or whatever takes its place). For four-dimensionalists then, including the contributions from future events does not pose such a great ontological problem as they already concede that these future events exist in exactly the same way as past and present events. Indeed, according to this view, explaining the

spectral creature's perception would be somewhat *less* problematic than explaining why we so strongly perceive the inescapable constancy of time's passage. For, the alien's temporal access would not be limited in this way and would therefore be commensurate with the equal status given to all times by four-dimensionalists. In such a situation, it seems unlikely that we could win the ontological argument in favour of spacetime, as the kind of operation in changing from the spatiotemporal picture to the spectral one becomes analogous to the change of reference frame in special relativity.⁹

Such a solution, however, cannot be accommodated by a second school of thought, albeit one currently rather less fashionable than four-dimensionalism, which states that the apparent passage of time is real. The class of future events *is* ontologically distinct from that of past events and our inability to access the future is not due to some limitation of our own perceptive abilities but directly results from the properties of time itself, the absence of backward-causation, etc. Is the need to include contributions from all times reconcilable with this theory about time? While not so readily accommodated as in four dimensionalism, there nonetheless seem to be two possible ways our creature could construct a spectral representation based on ω rather than t , while still admitting (in our language) a genuine flow of time. Because we now take the passage of time as real, we cannot give the creature access to all times, but must limit it at most to the temporal interval of its past world-line. Assuming this is of the order of several decades (rather than allowing t to span eternity) will not greatly affect what follows as only the perception of systems which evolve with characteristic times longer than this will be affected. But we humans would not directly perceive processes that take 10 years before a change is perceptible as temporally evolving either, so the parallel limit on the creature's perceptive abilities seems fair.

In the first (and least interesting) approach which takes time's passage as real, our creature is much like the ultra low-frequency spectrum analyser introduced in Section 4. The period of the alien's perceptual sampling rate is precisely its own lifetime so that it spends its entire life gathering data to construct the one, single perception event of its existence. It becomes "aware" for an instant just before it dies and literally sees its entire life "flash before its eyes", but as a spectral distribution rather than a temporal sequence. The creature would have no awareness of time, though of course all the time-dependent phenomena we would perceive during an identical lifetime will be encoded in the creature's frequency spectrum. This scenario is rather unsatisfying however, because it begs the question in favour of the temporal picture being fundamental. For, there is a feature of time, namely its passage, which we as temporally oriented creatures are aware of, but

which is not reflected in the spectral creature's world view. In addition, the spectral creature would presumably not be able to interact with the world around it in a conscious way until its "awakening" just before its death, making its survival a highly unlikely happenstance.

In the second, more interesting attempt to endow our creature with a spectral world-view consistent with an ontological commitment to the passage of time, there is a dynamic aspect to the frequency spectrum that parallels the flow of time. Rather than remaining an unconscious data gatherer for almost its entire life, this specimen from our spectral creature zoo is constantly perceiving. It represents its life "to date" by the frequency spectrum corresponding to what we would think of as the temporal sequence of events comprising its existence. But as it gets older, i.e. as the time interval we must integrate over gets larger, the lower bound on the frequency spectrum must get lower if all the accruing experience is to be cognitively accommodated. So in this creature's world-view, the time-increasing sequence is replaced by a frequency-decreasing sequence. To illustrate this crudely, suppose we think of a pair of perceptions which we loosely tag as being about 3 years and 30 years (or about 10^8 and 10^9 seconds) after the first thing we remember. In the creature's way of seeing the world, there would be analogous perception-states, one of which required a lower bound frequency of 10^{-8} Hz and another which required a lower bound frequency of 10^{-9} Hz (such fine distinctions would require enormous sensitivity, so the creature's actual representation may be something like $-\log(\omega)$, rather than simply ω . But that is a minor detail). Again, we would be hard pressed to convince such a creature that our naive ontology is correct, rather than its own. Indeed, a spectral-Kant may well have written (or whatever substitutes for writing) that "Wavenumber is a necessary *a priori* representation, which underlies all outer intuitions . . ." and a pair of spectral philosophers might attempt to put forward the apparently ludicrous proposition that there could exist a coherent conceptual scheme where space and time are more fundamental than wavenumber and frequency . . .

6. CONCLUSION

Irrespective of Kant's assertion that our cognitive framework necessarily depends on space and time, there exist formal methods where this dependence can be transformed out of the mathematical representations of physical systems. A particularly simple but physically important example of such a transformation is the Fourier transform. As shown in Equation (1), spacetime dependent functions are transformed to functions that are

instead dependent on wavenumber and frequency. The mapping is both invertible and hence works equally well in both directions, and is very similar in form to its inverse, therefore not privileging either direction of the transformation. Given that such an alternative mode of description exists that appears to be equivalent to a spatiotemporal one, we naturally wonder whether there exists a similar mode of perception. Again, the widespread utility of the Fourier transform in physics and engineering highlights that there already exist both artificial and biological sensors that are modelled by this analytical tool. The fact that our primary sense organ for constructing the spatial world view, namely the eye, seems not to primarily function as a Fourier transformer can be imagined to be at least partially responsible for the development of an incorrigible reliance on spacetime, but also as a somewhat accidental development.

Our thesis here regarding the consequences of this for spacetime should not be taken as reductionist, nor as eliminativist. We are not claiming either that spacetime is not fundamental because there is something else that is more fundamental, or that spacetime can be dispensed with altogether in favour of something else. Rather, if one wishes to draw ontological conclusions, our argument is most closely akin to a notion that must, at first, have seemed equally challenging: that of the complementarity principle of quantum mechanics popularly known as “wave-particle duality”. One does not wish to state either that the primitive constituents of matter are “really” waves or that they are “really” particles, but that these two notions provide limiting cases on which idealised models can be based. Depending on the details of the experiments one does, either of the two extremes can be highlighted.

Based on the above arguments, we claim that a similar complementarity holds between the spatiotemporal and spectral views of the physical world. The key difference between this and wave-particle duality is that, while we can perform quantum experiments that highlight either the wave or particle idealisations of matter, the physics of our sensory system is such that the “experiments” it performs on the world always conform to the spatio-temporal idealisation and the cognitive capacities we have evolved have been shaped by this. We suggest that this constraint be recognised as arbitrary and accidental, at least to the extent that our inhabiting an environment where such representations are efficient is a contingent consequence of our evolutionary circumstances. Furthermore, by admitting the possibility of sensory perception based on physical mechanisms that detect wave properties, we can construct an alternative representation of the world that is commensurable with the spatiotemporal, but radically different from it. Consequently, as Kant’s account of the role space and time

play in our perception implicitly suggests, the spatiotemporal framework for representing the world is, indeed, not unique.

NOTES

* The first author wishes to express his deep gratitude to Graeme Rhook, who after many stimulating and vigorous arguments, helped shape several of the ideas in this paper into their final form. The second author is grateful to La Trobe University, Dept. of Physics, and in particular to his host Professor P. L. Dyson, and the CRA (Conzinc Riotinto Australia) consortium for providing support for this research. Thanks are due to friends and colleagues willing to listen and contributing to the subject by their comments. In particular, the second author wishes to thank members of DSTO (Defense Science and Technology Organization) in Salisbury (Adelaide), South Australia, in particular John A. Bennett (also of the Department of Electrical Engineering, Monash University, Clayton (Melbourne), Victoria, Australia), Stuart Anderson, and Chris Coleman, for lively arguments and drawing our attention to various references. The second author is also grateful to colleagues in the Ben Gurion University, Beer Sheva, Israel for comments, especially Tsilil Censor of the Department of Philosophy, for listening, arguing and providing references.

¹ Perhaps the use of “momentum space” representations in quantum theory is even more graphic an illustration of our idea. But this would introduce a level of mathematical complexity that adds nothing of any substance to our discussion, and might lead one to think that our arguments rely on “mysterious” quantum properties. As we wish to do neither of these things, all our physical examples will be couched completely in terms of familiar, classical systems.

² One who has stared at a colour TV screen or computer monitor up close, or has programmed computer graphics may object at this point and claim that white requires only high intensities of red, green and blue. This, however, is a peculiarity of our visual system as the receptors in our eyes that enable us to distinguish colours are only sensitive to these three colours (although their excitation spectra are far from monochromatic). Hence, if our retinas are strongly stimulated by red, green and blue light, we undergo the subjective experience of seeing white light, even though a reasonable fraction of the frequency spectrum (particularly between green and blue to which our eyes are not particularly sensitive) may be quite diminished.

³ For further details on the mathematics, consult almost any textbook on functional analysis or on the theory of complex functions, e.g. Fisher (1990), pp. 318–346.

⁴ It might be argued that in order to define v , one must refer to spatiotemporal quantities because $v = dx/dt$. But this is not the case, as simple algebra will demonstrate that the above relation when applied to the plane wave with constant phase $\omega t - kx$ also implies $v = \omega/k$, i.e., the phase velocity. Alternatively, one could express v as $v = dw/dk$, i.e. the group velocity, which is more similar in form to the spacetime expression and represents something more akin to what we would mean by the velocity of a wave propagating through space.

⁵ Although proprioception and kinaesthetics must also play some part, we shall not consider them explicitly here because the spatial information they give us cannot be as precisely expressed as is the case for visual images. Also, the fact that our bodies are solid, irregular, and have precise enough boundaries makes the relative position of our limbs a meaningful notion and a useful property to sense. But consider whether this would be true for some

kind of highly symmetric life-form (say a perfect sphere) or ones that exist as intermingled and highly distributed liquid or gaseous forms in a similarly amorphous environment.

⁶ It might be mistakenly assumed that the “wavelength” properties being referred to here are simply the wavelengths of the light entering the creature’s eyes. While this is true to an extent, the properties we are interested in are the *spatial wavelengths* of the objects in the visual field, i.e. the way that these objects can be represented as superpositions of oscillatory patterns (just as they can be represented as sets of points – the mode familiar to us).

⁷ We should take care at this point and point out that strictly, the ear is only sensitive to the *intensity* of the component modes. It cannot generally detect phase.

⁸ There are, of course, blind people who do not have visual sensations. However, they have brains which evolved from creatures that did and so we assume that their representation of the world is based on the same spatiotemporal framework as their sighted counterparts.

⁹ As we mentioned above, it is relatively easy to construct a technical argument showing that the Fourier transform in functional analysis is the equivalent of a basis transformation in vector analysis. However that is not the subject of this paper.

REFERENCES

- Abraham, M.: 1904, ‘Zur Theorie der Strahlung und des Strahlungsdruckes’, *Annalen der Physik* **14**, 236–287.
- Censor, D.: 1991a, ‘“Waves”, “Objects” and Special Relativity’, *Journal of Electromagnetic Waves and Applications* **5**, 1365–1391.
- Censor, D.: 1991b, ‘Application-Oriented Relativistic Electrodynamics’, in J. A. Kong (ed.), *PIER 4: Progress in Electromagnetics Research*, Elsevier, New York.
- Dawkins, R.: 1991, *The Blind Watchmaker*, Penguin Books, London.
- Doppler, C.: 1842, ‘Über das farbige Licht der Doppelsterne und einiger anderer Gestirne Des Himmels’, *Abhandlungen der Königlich Böhmisches Gessellschaft der Wissenschaften* **1**, 467–482.
- Einstein, A.: 1923, ‘On the Electrodynamics of Moving Bodies’, in Lorentz, Einstein, Minkowski and Weyl (eds.), *The Principle of Relativity*, English Translation by W. Perrett and G. B. Jeffery. Originally published as ‘Zur Elektrodynamik bewegter Körper’, *Analen der Physik* **17** (1905), Dover, New York.
- Fisher, S.: 1990, *Complex Variables*, 2nd. Edition, Wadsworth & Brooks/Cole, Pacific Grove, California.
- Gill, T. P.: 1965, *The Doppler Effect*, Academic Press, New York.
- Jackson, J.: 1975, *Classical Electrodynamics*, 2nd. Edition, Wiley, New York.
- Kant, I.: 1933, *Critique of Pure Reason* (original 1781), Translated by Norman Kemp Smith, second edition, Macmillan, London.
- Nagel, T.: 1979, *Mortal Questions*, Cambridge University Press, New York.
- Pauli, W.: 1958, *Theory of Relativity*, Pergamon Press, New York.
- Poincaré, H.: 1905, *Science and Hypothesis*, Dover edition published 1952, New York.
- Reichenbach, H.: 1958, *The Philosophy of Space and Time*, Translated by M. Reichenbach and J. Freund, originally published as *Philosophie der Raum-Zeit-Lehre* (1927), Dover, New York.
- Strawson, P. F.: 1959, *Individuals: An Essay in Descriptive Metaphysics*, Methuen & Co. Ltd., London.

Mark Zangari
Philosophy Department
La Trobe University
Bundoora
Victoria 3083
Australia
M.Zangari@latrobe.edu.au

Dan Censor
Department of Electrical
and Computer Engineering
Ben Gurion University of the Negev
Beer Sheva
Israel
Censor@bgu.ac.il