Purposeful Choice and Point-of-View
A Generalized Quantum Approach

Dominic Widdows

Serendipity

Presented at 7th Quantum Interaction Conference, Leicester, UK, 2013.

Abstract. This paper presents a generalized quantum model for describing purposes or goals of individual agents, and the way choices can be made that enable these goals to be achieved. The underlying model is a semantic vector space model, which is turned into a purposeful choice model by labelling some axes as objectives, and describing choices as transformations on the vector spaces that enable agents in the model to set these objective axes in sight.

We introduce this framework using a simplified example model of a dog trying to get food. Many parts of what has become the standard generalized quantum toolkit become apparent in this model, including learning, superposition, the importance of the metric used for normalization, classification, and a generalized uncertainty principle. The incorporation of purpose or goal into semantic vectors models also enables the analysis of traditional areas that are relatively new to artificial intelligence, including rhetoric, political science, and some of the philosophical questions touched by quantum theorists.

1 Introduction

This paper presents a vector model for the everyday notions of 'purpose' and 'choice'. While the initial example model requires as background only the basics of Cartesian coordinate geometry, as it progresses, the work builds on progress in quantum interaction and generalized quantum structures, which have in recent years been used with success to address many classic problems. The basic approach in this research program involves representing information systems using vector space models, and as such has been applied to information discovery and retrieval [1, 2], cognition and decision-making [3], economics [4, Ch. 10] and organizational dynamics [5]. The 'quantum' qualities of these systems evaluated in the literature to date include non-locality of logical connectives in information retrieval [6, Ch 7], non-commutativity of observables in psychological tests [3], the violation of the Bell inequalities in concept combinations [7, 8], and entanglement in concept combinations [9, 1].

In spite of the success of vector models in large-scale practical tasks such as the creation of information retrieval systems, the notions of purpose and choice in such models is comparatively unexplored, sometimes being ignored, and sometimes being entirely denied as a valid ingredient of the model in the first place. Scientific works where purpose is largely ignored include studies of the parallels between information retrieval and quantum theory (see e.g., [10], where the epistemological and ontological status of items in the system is considered, but the motivations for creating or using the system are not), and distributional models of concepts and their semantics derived from natural language corpora (see e.g., [11, 12]). Of course, the relationship of items in a retrieval system to one another and to objects in the world, and the distribution of terms and topics in a corpus, are important and valuable areas of study in understanding language and meaning: however, they do not attempt to explain anything about what authors are trying to accomplish by writing documents, or what users or a retrieval system are trying to accomplish by issuing search queries. Scientific examples where purpose is denied as a valid ingredient of the model are much more general, and are a hallmark of many classical mechanical approached. In broad strokes, Francis Bacon's philosophy and the success of Newtonian mechanics in the 17th century led to a broad consensus in the 18th century that the only notion of cause that can be discussed scientifically is 'efficient cause' or cause in the mechanical sense: one of the few things Hume (a great empiricist) and Kant (a great rationalist) agreed upon was the notion that causes must precede their effects in time. This overrode Aristotle's much older analysis in which included 'final cause' or purpose among the natural causes of things (see in particular *Physics* Book II Ch. 3; for modern consequences see [13, Ch. 1]).

Whether or not the apparent notion of future purpose can be explained in terms of temporally prior mechanical causes (for example, by a generalization of field theories in which the notion of potential is explained in terms of force-carrying particles), it is noticeable that classic models motivated by cause-precedes-effect determinism have been found wanting in many fields in which the systems under consideration are too complex or subtle to be described as closed mechanical systems evolving predictably [14], and the reader will observe that many of the successes of generalized quantum approaches cited above are precisely in fields that are not (yet) amenable to mechanistic prediction. In simpler terms, as soon as we consider systems involving living things, especially people, we see that purpose and choice are fundamental factors in any thorough explanation. These cannot (yet) be explained in terms of more mechanical primitives, but cannot be neglected if effective scientific models of such systems are to be discovered. This is by no means a purely abstract exercise: appropriate models for understanding the goals and choices of authors and readers could (for example) enable engineers to build better search engines.

It should be noted that models for purpose and choice are not absent from the scientific literature: one of the most famous approaches is the use of 'belief, desire, intention networks', which [15] have been particularly influential in modelling agency in artificial intelligence [16]. Generalized quantum methods are potentially a complementary innovation to such discrete network models, because the continuous vector representation automatically enables robust or inexact inference, in ways that are naturally amenable to learning from experience [17, 18]. Decision-making is also by now deep-rooted in the quantum interaction community (see particularly [3] and associated works). Here we note that most of the decision-making situations discussed in this literature are about deciding between possible information states or beliefs: so arguably, the innovative part of the purposeful choice model presented in this paper is that it applies vector representations to desires and intentions as well as beliefs. However, it is also our hope that this research area is by now mature enough that the contribution of this paper is not that it supersedes prior work, but that it simplifies, generalises and extends ideas that are already available.

With these goals in mind, this paper proceeds as follows. In Section 2, we introduce a first, extremely simple example that explains the behaviour of an agent (in this case, a family pet) in a model with one objective axis and two axes for expressing behavioural choices. The semantic space introduced in this model is similar to the distributional vector models used widely in information retrieval, computational linguistics, and cognitive science, but unlike many semantic models in these fields, the purposeful choice model presented here distinguishes 'ends' and 'means' directions.

In spite of its simplicity, this model is enough to motivate definitions for several important cognitive processes, including learning and classification: some of these developments are discussed in Section 3. Many further topics and developments are suggested by this discussion, but due to space constraints, they cannot be included in this paper. Section 4 outlines some of these topics. They include the modelling of rhetoric, applications to political theory, and the relationship between purposeful choice models and some standard areas of discussion in the philosophy of quantum mechanics.

2 First Example: A Dog's Life

"Look cute, get fed!" may be the motto of fortunate, well-kept pet dogs throughout many parts of the world. Many readers who own dogs are probably well aware of this trait: for those who are not, it is sufficient to note that:

1. Most dogs (especially those rescued from situations of hunger) are tirelessly devoted to the purpose of getting food.
2. Pet dogs devote themselves to this purpose by seeking out humans who might give them food, and doing their best to look cute, cuddly, hungry, pitiful, and attractive to humans as best they can.

In the wild, these behavioural traits are not especially useful compared with the basic hunting skills of (say) being able to run fast to catch prey. However, pet dogs have successfully transformed their strategy for getting food from running fast to looking cute. Experiments with the domestication of the silver fox, a closely related species to the dog, have demonstrated that profound behavioural changes can take place within a matter of a few decades, or 30 to 35 generations, leaving the animals "eager to please and unmistakably domesticated" [19]. By the same token, many tame dogs can make use of their wilder traits at a moment's notice: ill-trained or uncontrolled sheepdogs may chase herds, and most terriers and hounds will kill any small furry creature given the opportunity. We may sum up by noting that all dogs wish to get food, different strategies are available (even to an individual dog), and the choice between these strategies is sometimes made quite quickly and fluidly.

A *purposeful choice model* for this (obviously simplified) description of a pet dog's objective and behaviours is represented in Figure 1. The model uses a 3-dimensional vector space, whose axes are labelled *Get Food*, *Run Fast*, and *Look Cute*. The *Get Food* axis is represented by a thicker line because it represents a purpose, otherwise described as a goal, end, or objective. Such an axes will be called an *objective axis*. The other two axes, *Run Fast* and *Look Cute*, represent different possible behaviours that a dog may choose between to achieve its goal.

The choice between behaviours is now modelled as a change of *point-of-view*. The use of this term in the model is just a formalization of its normal conversational use: a point-of-view is a place from which the concepts around are observed. So an agent adopts a point-of-view in the model. The wild dog adopts a point-of-view which keeps the *Get Fed* objective axis in sight and approximately aligns this axis with the *Run Fast* choice. The tame dog instead adopts a point-of-view which approximately aligns the *Get Fed* objective axis with the *Look Cute* axis.

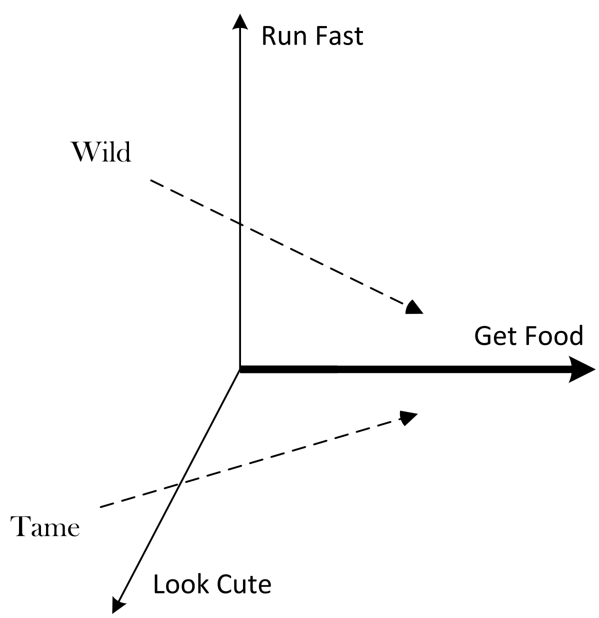

Fig. 1. A purposeful choice model for a dog in three dimensions

The views from the point-of-view of the wild and the tame dog are shown in Figure 2. For the wild dog, the objective axis *Get Fed* is aligned with the *Run Fast* axis, whereas for the tame dog, the *Get Fed* axis is aligned with the *Look Cute* axis. The key point to see is that by adopting a different point-of-view, the relationships between different axes change, and that this alignment can be made very deliberately to align behaviours with objectives.

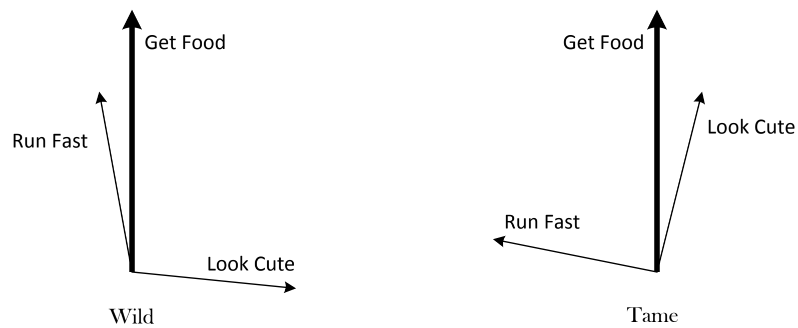

Fig. 2. The dog model again, from the points-of-view of a wild dog (left) and a tame dog (right)

The rest of this paper, one way or another, will be devoted to fleshing out this basic model, and explaining how more sophisticated (and to the research community, more familiar) structures arise in this framework.

The reader should note before progressing further that we have made no claims or assumptions of orthogonality, or assumed any particular metric function in the purposeful choice model. Many well-kept pet dogs will agree that looking cute and running fast are not necessarily orthogonal, and indeed, the different strategies for accomplishing a particular objective are rarely entirely unrelated to one another. To model a purposeful choice as a change in point-of-view gives considerable freedom here: even if axes are orthogonal to each other in the underlying model given a particular metric, they will not usually appear orthogonal to each other once a point-of-view is selected.

The investigation of similarity methods with respect to a point-of-view has already been introduced in [20], and again, the main single innovation introduced in this paper is the use of these ideas to model behaviour directed towards particular purposes or objective. The transformation in similarity measurements resulting from a change in point-of-view could also be modelled using a change in the metric function on the vector space, or a rescaling of some of the axes. This is also suggested in the cognitive science literature: changing the weights assigned to different axes can change the way items are classified in experiments [21].

From a historical point-of-view, we note that orthogonal coordinate systems are not assumed in Descartes' pioneering work on analytic geometry [22] in the 1600's, though it is explicitly discussed in Grassmann's extension of the theory to higher dimensions in the 1800's [23].

Another point to note is that, whatever the relationship between the *Run Fast* and *Look Cute* axes, most dogs use a superposition of these states anyway. Wolves may be single-minded to running fast, shih tzus may be single-minded to looking cute, but most real dogs have a foot in each camp. As will be discussed in the next section, vector models are particularly well-suited to representing hybrid strategies of this sort.

Of course, this description is purely a mathematical model, and as such is a simplification and abstraction. We are not attempting to model the physiological or neuronal patterns and changes involved in transitioning from one strategy to another (as discussed in the case of canids in [19] and mentioned briefly from a cognitive point-of-view in [17]). Just as vector models for information retrieval do not describe the physical formats of documents (typefaces, character encodings, etc.), our vector model for purposeful choice is so far independent of its physical manifestation.

This concludes our initial presentation of the purposeful choice model, using the simplest possible nontrivial number of dimensions. The key parts to emphasize are that:

- Purposeful choice models are semantic vector models where some of the axes are marked as goals or objective axes.
- Other axes can be brought into line with these objective axes using a suitable transformation of point-of-view.

3 Common Structures in the Purposeful Choice Model

This section develops the ideas of the purposeful choice model introduced above. This serves two principal purposes. Firstly, it demonstrates how several well-established techniques can be incorporated and described in terms of these models. Secondly, where it presents itself, we take the opportunity to compare classic and generalized quantum models: in some cases they are similar, but it some cases they lead to strikingly different paths.

3.1 Learning in Purposeful Choice Models

Vector models are particularly well-suited to learning, and this is one of the key reasons they have been a key model in information retrieval [24, 2] and have become so successful in statistical machine learning [25]. This is easy to explain, at least anecdotally, referring back to the dog model of Figure 1. Suppose that a piece of food is available through hunting: then a dog who gets fed in this way learns that hunting characteristics are useful. Such an example is modelled as a point somewhere close to the plane spanned by the *Run Fast* and *Get Food* axes. If the dog gets food in this way, it ends up satisfying an objective, and in so doing, the dog's point-of-view is updated to align these axes more closely. Alternatively, if food is available through begging, this may be modelled as a point somewhere near to the *Look Cute* and *Get Food* axes, and a dog that successfully fills its belly through begging will have its point-of-view updated to align these axes. (Which of several possible update functions is used is not discussed here, suffice it to say that several are available [26, 25].)

Such flexibility to combine learning with action would in itself not be remarkable, were it not so lacking in many classic models. Consider, for example, Quine's now famous example of the problem of deducing whether the word *gavagai* coupled with the stimulus of a rabbit-sighting, corresponds to the set of rabbits, or to (say) edible animals [27]. Given the practical advances of empiricism in artificial intelligence in the intervening decades, many researchers today would disagree with the conclusion that language-learning cannot be explained logically, but would instead argue that to model language-learning, one should use a more appropriate logic. (See [28, 18, 29] for more details on this point: it is also appropriate to note that George Boole, the inventor of so-called classical logic, intended 'The Laws of Thought' to be used for deduction, and apparently never intended to apply them to learning [30].)

3.2 Objective Axes and Objective Functions

The most typical way to compare directions in a vector model would be to use cosine similarity (that is, the similarity between two vectors is measured using the cosine of the angle between these vectors [6, Ch 5]. Cosine similarity with an objective axis behaves as a simple *objective function* in the classic sense of mathematical optimization, as used in economics, logistics, management science, etc. At this level of generality, with a single objective axis and an entirely specified point-of-view, there appears to be no difference between a classic and a generalized quantum model. What is perhaps more surprising is that in mathematical optimization, classic models tend to be continuous, whereas in logical semantics, classic models are discrete. This mixture is partly informed by the observation that in classical mechanics, the set of states is continuous but the logic for inference is discrete (Boolean), whereas in quantum mechanics, the set of states is discrete but the logic for inference (the set of projectors onto vector subspaces [31]) is continuous. For more details on this point, see [32]. This serves as a reminder that definitions of a 'classic' or 'classical' model vary even more than definitions of 'quantum' or 'generalized quantum' models.

3.3 Superposition, or Hybrid Strategies

It has already been noted that most real dogs are not devoted exclusively to running fast or looking cute as a means to get food, but easily combine both strategies. Suppose that x is the amount of attention devoted to scenarios where looking cute is helpful, and y is the amount of attention devoted to scenarios where running fast is helpful. In a classic probabilistic model, the coordinates must add to unity, so $x + y = 1$. For (say) a wild dog accustomed purely to running fast (the strategy where $y = 1$ and $x = 0$, giving any attention to looking cute therefore immediately detracts from the attention devoted to running fast. By contrast, one key different in a generalized quantum model is that the *squares* of the coordinates must add to unity, so that $x^2 + y^2 = 1$. An immediate consequence of this is that beginning to pay attention to another axis introduces *no loss* [1] to the axes that are already preferred. These two frameworks are depicted in Figure 3: in mathematical terms, the vectors in the classic model are normalized using a Manhattan

[1] The claim that there is no loss at all to existing priorities when a new axis is considered is strictly true only in the continuous limit. For example, using the standard polar-coordinates parametrization where $x = \cos(\theta)$ and $y = \sin(\theta)$, when $\theta = 0$, $x = 1$, $y = 0$, $\frac{dy}{d\theta} = 1$ and $\frac{dx}{d\theta} = 0$. In practice, we assume that all models will be quantized, and so to make an actual change, there will be some *very small* cost. The issues involved in quantizing vector models of information and cognition is not a focus of this paper: we note briefly that this example implies that it is advantageous for the smallest 'representational quantum' to be small.

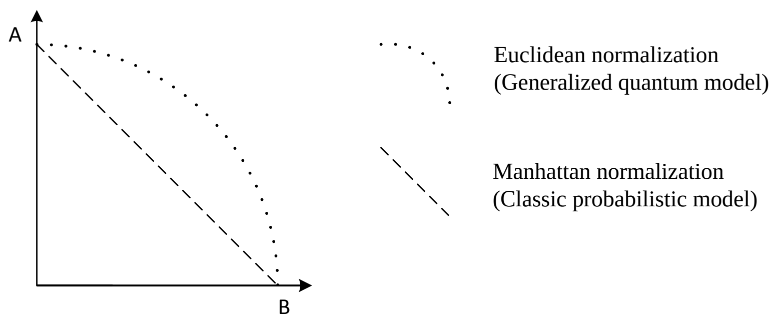

Fig. 3. The set of normalized states between two axes in classic and generalized quantum models

metric, whereas in a generalized quantum model, the vectors are normalized using a Euclidean metric (see [6, Ch 4]).

The upshot of this is that in generalized quantum models, there is very little cost involved in departing from a pure state to a somewhat mixed state. The result is a dog whose ideal strategy is modelled as a superposition of the two pure strategies. This is a difficult idea to incorporate into classic models: in classic models, we may model a dog who sometimes tries one thing, and sometimes tries another, but not a dog who is simultaneously trying both. However, such a description has become standard in the quantum modelling literature, and has been used to accurately predict experimental observations about how people behave in situations where many potential outcomes are simultaneously possible: see in particular the work of Busemeyer et al ([33],[3, Ch. 9]) in which shows that difficult situations in psychology that have been traditionally regarded as paradoxes can be resolved using these methods.

We must of course choose words carefully in this case: clearly the dog cannot normally be pursuing prey and begging from a human at the same time, and we make no suggestion that the benefits of Euclidean normalization of a point-of-view vector can somehow enable an agent to outwit conservation of energy. The claim here is that agents can easily *consider* several strategies at once without diluting the attention devoted to any one of these strategies as much as may be expected in a classic model. A more familiar version of this principle arises when it is translated to the information retrieval literature: here we may remark that in vector models for information retrieval, if a document is about two topics, its relevance to each of the individual topics as measured by cosine similarity will be $\frac{\sqrt{2}}{2}$ instead of just 0.5. Furthermore, this property becomes even more pronounced in higher dimensions: in very high dimensions, many many vectors can be superposed without losing the identity of any of the original summands [17]. [2]

[2] The difference between normalized coordinates of evenly-balanced vectors using Manhattan and Euclidean metrics is greatest in dimension 4. The proof is elementary, and consists of

The immediate consequence in the simple purposeful choice model of Figure 1 is that it costs very little in terms of cognitive attention for a wild dog devoted to hunting to consider begging as a once-in-a-while alternative. This potential for an individual to break free from the rest of the population and find a new strategy has already been modelled successfully in quantum-inspired models (see [34]), so we hope that the purposeful choice model presented here contributes to this strand of research.

3.4 Classification — Fight or Flight?

The purposeful choice model of Figure 1 is oversimplified in many ways, one of which is that the dog only has one goal, *Get Food*. Obviously a real dog has several other objectives, including *Avoid Injury*. For example, in the case of a dog trying out a *Look Cute* begging strategy for the first time, there is a tradeoff between the possibilities that a human offering a piece of food may give the piece of food, or use it as bait to capture or injure the dog, and there are good (if anecdotal) reasons to believe that the first dogs to become domesticated were the first to overcome this fear.

Consider also the case of a dog who has successfully obtained a piece of food, but is challenged by an antagonist before the food is eaten. There are two choices: to stay and safeguard the food at the risk of injury, or to run away, at the risk of hunger. This is commonly known as the *flight or fight* decision.

It is extremely easy to begin to add such tradeoffs to a purposeful choice model. We would simply add a *fourth* axis to Figure 1, marked *Suffer Injury*, and with some label to denote that fact that it is a *negative* objective axis: that is, an axis that the agent will try to *avoid* aligning with behaviours in the purposeful choice model. (Due to lack of space and our inability to draw in four dimensions, we have not included a diagram of such a system.) Then, for each situation, the agent judges the extent to which *Get Food* and *Suffer Injury* are likely outcomes, and, based on some decision boundary, will choose either fight or flight (alternatively, beg from human or retreat) accordingly. Again, several algorithmic strategies for learning such decision boundaries are available in the machine learning literature [26, 25].

3.5 A Generalized Uncertainty Principle

An important consequence of choosing *flight* in a *flight or fight* decision is that the agent is unable to observe the outcome of the other decision. This is very obvious in everyday situations, and leads to natural sayings such as “There’s only one way to find out!”, “If you don’t try it you’ll never know!”, etc.

More generally, in navigating a purposeful choice model, an agent will be aware that each time a choice is made, this affects which observations will be distorted or become completely unavailable at subsequent stages. This idea follows the work of Busemeyer and Bruza on the effects of ordering on attitude (see [3, Ch 3] and related work). In particular, the effect of making a *flight* choice may be modelled as a projection

finding $x \in [0,1]$ such that $f(x) = \sqrt{x} - x$ is maximized, so $f'(x) = \frac{x^{-\frac{1}{2}}}{2} - 1 = 0$, implying $x^{\frac{1}{2}} = \frac{1}{2}$ and so $x = \frac{1}{4}$. We are not sure if this number has any special significance.

orthogonal to the *Suffer Injury* axis, which, while it guarantees that the agent will avoid injury, also results in information being lost.

4 Further Work and Related Areas: A Grab-bag of Ideas

Many traditional ideas can be defined and described in purposeful choice models. We use the remaining space in this paper to outline some of these in a preliminary fashion.

4.1 Persuasion or Rhetoric

It is well-known that information is often presented in a way designed to persuade or influence the point-of-view of others. Traditionally studied as rhetoric, the scientific discussion of this hugely important process has been largely neglected in computational linguistics and information retrieval.

In a purposeful choice model, persuasion or rhetoric can be defined as the presentation of information in a way designed to influence the agent's point-of-view. Given the approach to training and classification outlined above, it is clearly possible to arrange data so that some points-of-view become reinforced, and others become less likely or unobtainable.

4.2 Application to Political and Organizational Theory

Generalized models have already been applied to political theory (see [35]) and organizational theory (see [36] and related work). As with work on quantum approaches to cognition and decision, the focus of this work is largely on describing how agents make decisions: a further step would be to model the ways other agents act in order to influence these decisions. Such influencing actions can be decribed in purposeful choice models as:

- A careful choice of issues by some author to design an appropriate classification boundary (e.g., in politics, a bill before the legislature is designed to accomplish as many of the author's desired goals, while maximizing the bill's chances of being voted into law).
- A careful choice of rhetoric designed to bring others to a point-of-view from which they are likely to agree with the author.

For example, President Lincoln's 1861 State of the Union Address makes an admirable case study of the use of rhetoric to align many points-of-view towards a common goal. Further work would be to model parts of this speech and its goals explicitly using a purposeful choice model.

4.3 The Purposeful Choice Model and Quantum Mechanics: Some Philosophical Conclusions

The notion of choice and decision is of course intimately connected to the idea of will in the sense of freedom of the will, a topic that is often associated with quantum theory, due largely to the probabilistic nature of quantum mechanical results. The notion of point-of-view is also relevant in quantum mechanics because in quantum mechanics, the observer is usually considered as part of the system, though the meaning and implications of this broad statement remain much-discussed [13].

It may be thought that a purposeful choice model implies an assumption of free-will at the expense of determinism. This is not necessarily the case. What *is* necessarily the case is that purposeful choice models agree with the basic Aristotelean principle that *final cause* is a valid and valuable kind of causation or explanation when studying natural processes ('natural' including human behaviour for these purposes). That is, a dedicated determinist may postulate that human behaviour including the notion of purpose could in principle be reduced to efficient or mechanical cause (just as the apparent 'goal' of an electron and a proton to be near each other can be explained mechanically by the exchange of photons). The problem with this approach is that it remains a very incomplete postulate. Social and information sciences have needed models that explain more of the phenomena observed in these disciplines, and this is what originally motivated many researchers to turn to generalized quantum models.

We cannot currently explain human (or even canine!) behaviour without the notion of purpose and choice: and even if it could be demonstrated in the end that such notions can be reduced to mechanical or efficient cause, models that successfully incorporate purpose into informatics would be a necessary stepping-stone. Thus, whether we are completing classical science or starting a new generalized science, the notion of purpose will be a key part of the explanation, and we suggest that the purposeful choice models introduced in this paper can play a valuable and practical role in this project.

References

1. T. Cohen, D. Widdows, L. de Vine, R. Schvaneveldt, and T. C. Rindflesch, "Many paths lead to discovery: Analogical retrieval of cancer therapies," [37].
2. K. van Rijsbergen, *The Geometry of Information Retrieval*. Cambridge University Press, 2004.
3. J. Busemeyer and P. D. Bruza, *Quantum models of cognition and decision*. Cambridge: Cambridge University Press, July 2012.
4. A. Khrennikov, *Ubiquitous Quantum Structure: From Psychology to Finance*. Springer, 2010.
5. W. F. Lawless, M. Bergman, J. Lou, N. N. Kriegel, and N. Feltovich, "A quantum metric of organizational performance: Terrorism and counterterrorism," *Computational and Mathematical Organization Theory*, vol. 13, pp. 241–281, 2007.
6. D. Widdows, *Geometry and Meaning*. CSLI Publications, 2004.
7. D. Aerts, S. Aerts, J. Broekaert, and L. Gabora, "The violation of bell inequalities in the macroworld," *Foundations of Physics*, vol. 30, pp. 1387–1414, 2000.
8. D. Aerts, S. Aerts, and L. Gabora, "Experimental evidence for quantum structure in cognition," (Saarbrücken, Germany), 2009.
9. D. Galea, P. Bruza, K. Kitto, D. L. Nelson, and C. McEvoy, "Modelling the activation of words in human memory: The spreading activation, spooky-activation-at-a-distance and the entanglement models compared," in *Proceedings of the Fifth International Symposium on Quantum Interaction*, pp. 149–160, 2011.
10. S. Arafat, "Senses in which quantum theory is an analogy for information retrieval," [38].
11. T. Cohen and D. Widdows, "Empirical distributional semantics: Methods and biomedical applications," *Journal of Biomedical Informatics*, vol. 42, no. 2, p. 390, 2009.

12. D. M. Blei, “Probabilistic topic models,” *Commun. ACM*, vol. 55, pp. 77–84, Apr. 2012.
13. D. Bohm, *Wholeness and the Implicate Order*. Routledge Classics, republished 2002, Routledge, 1980.
14. D. Widdows and P. Bruza, “Quantum information dynamics and open world science,” in *Proceedings of the First International Symposium on Quantum Interaction*, (Stanford, California), 2007.
15. M. E. Bratman, “Intention, plans, and practical reason,” 1999.
16. M. Georgeff, B. Pell, M. Pollack, M. Tambe, and M. Wooldridge, “The belief-desire-intention model of agency,” in *Intelligent Agents V: Agents Theories, Architectures, and Languages*, pp. 1–10, Springer, 1999.
17. P. Kanerva, “Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors,” *Cognitive Computation*, vol. 1, no. 2, pp. 139–159, 2009.
18. P. D. Bruza, D. Widdows, and J. Woods, *A Quantum Logic of Down Below*, pp. 625–660. Elsevier, 2009.
19. L. Trut, “Early canid domestication: The farm-fox experiment foxes bred for tamability in a 40-year experiment exhibit remarkable transformations that suggest an interplay between behavioral genetics and development,” *American Scientist*, vol. 87, no. 2, pp. 160–169, 1999.
20. S. Aerts, K. Kitto, and L. Sitbon, “Similarity metrics within a point of view,” [38].
21. P. Gärdenfors, *Conceptual Spaces: The Geometry of Thought*. Bradford Books MIT Press, 2000.
22. R. Descartes, *The Geometry of René Descartes*. Dover, 1637. Dover edition, with facsimile of the original.
23. H. Grassmann, *Extension Theory*. History of Mathematics Sources, American Mathematical Society, London Mathematical Society, 1862. Translated by Lloyd C. Kannenberg (2000).
24. R. Baeza-Yates and B. Ribiero-Neto, *Modern Information Retrieval*. Addison Wesley / ACM Press, 1999.
25. T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning*. Springer Series in Statistics, 2001.
26. T. Mitchell, *Machine Learning*. McGraw-Hill, 1997.
27. W. V. O. Quine, *Word and object*. MIT Press (MA), 2013 (orig. 1960).
28. D. Gabbay and J. Woods, “The new logic,” *Logic Journal of IGPL*, vol. 9, no. 2, pp. 141–174, 2001.
29. D. Widdows and M. Higgins, “Geometric ordering of concepts, logical disjunction, learning by induction, and spatial indexing,” in *Compositional Connectionism in Cognitive Science*, (Washington, DC), AAAI Fall Symposium Series, October 2004.
30. G. Boole, *An Investigation of the Laws of Thought*. Macmillan, 1854. Dover edition, 1958.
31. G. Birkhoff and J. von Neumann, “The logic of quantum mechanics,” *Annals of Mathematics*, vol. 37, pp. 823–843, 1936.
32. V. S. Varadarajan, *Geometry of Quantum Theory*. Springer-Verlag, 1985.
33. J. R. Busemeyer, E. M. Pothos, R. Franco, and J. S. Trueblood, “A quantum theoretical explanation for probability judgment errors,” *Psychological Review*, vol. 118, no. 2, p. 193, 2011.
34. K. Kitto and F. Boschetti, “The quantum inspired modelling of changing attitudes and self-organising societies,” [37].
35. C. Zorn and C. E. S. Jr., “Pseudo-classical non-separability and mass politics in two-party systems,” [38].
36. W. F. Lawless and D. A. Sofge, “Social-psychological harmonic oscillators in the self-regulation of organizations and systems,” [37].
37. *Sixth International Symposium on Quantum Interaction*, (Paris, France), 2012.
38. *Fifth International Symposium on Quantum Interaction*, (Aberdeen, UK), 2011.