Meaning updating of density matrices

Bob Coecke and Konstantinos Meichanetzidis
Oxford University, Department of Computer Science
Cambridge Quantum Computing Ltd.

This is a three-page abstract of a paper with the same title, currently under review as an invited submission for a SEMSPACE special issue of the Journal of Applied Logics. The full version is `arXiv:1904.03478`.

Abstract

The DisCoCat model of natural language meaning assigns meaning to a sentence given: (i) the meanings of its words, and, (ii) its grammatical structure. The recently introduced DisCoCirc model extends this to text consisting of multiple sentences. While in DisCoCat all meanings are fixed, in DisCoCirc each sentence updates meanings of words. In this paper we explore different update mechanisms for DisCoCirc, in the case where meaning is encoded in density matrices—which come with several advantages as compared to vectors.

Our starting point are two non-commutative update mechanisms, borrowing one from quantum foundations research [5, 6], and the other one from [3, 7]. Unfortunately, neither of these satisfies any desirable algebraic properties, nor are internal to the meaning category. By passing to double density matrices [1, 10] we do get an elegant internal diagrammatic update mechanism.

We also show that (commutative) spiders can be cast as an instance of the update mechanism of [5, 6]. This result is of interest to quantum foundations, as it bridges the work in Categorical Quantum Mechanics (CQM) with that on conditional quantum states. Our work also underpins implementation of text-level Natural Language Processing (NLP) on quantum hardware, for which exponential space-gain and quadratic speed-up have previously been identified.

First we recast DisCoCirc as a theory of updating e.g.:

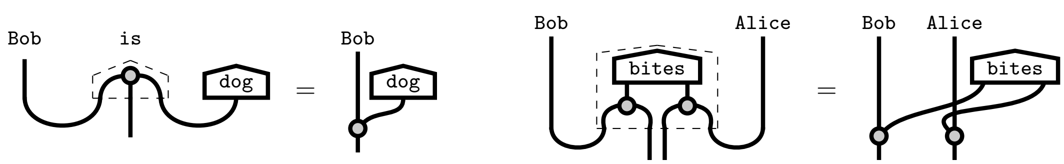

In this form, agents Alice and/or Bob are updated with `(being a) dog` and `(being in a) biting relationship` respectively. This idea applies to essentially all diagrammatic constructs, so:

What is the grey dot?

We consider this question for the specific case that of meanings as density matrices.

Following [9, 8] it is natural to think of propositions for natural language meaning like `(being a) dog` as a projector, just as in Birkhoff-von Neumann quantum logic [2]. Imposing a proposition on a density matrix is then realised as follows:

$$P \circ \rho \circ P \tag{1}$$

How does this generalise to imposing a density matrix? Two options have already been proposed:

ρ  $\sigma := \sum_i x_i \Big(P_i \circ \rho \circ P_i \Big)$ with $\sigma := \sum_i x_i P_i$ (2)

ρ  $\sigma := \left(\sum_i x_i P_i \right) \circ \rho \circ \left(\sum_j x_j P_j \right)$ with $\sigma := \sum_i x_i^2 P_i$ (3)

The first of these, which we call *fuzz*, was proposed in [3, 7] specifically for NLP. The second one, which we call *phaser*, was proposed in the form $\sqrt{\sigma}\,\rho\,\sqrt{\sigma}$ within the context of a quantum theory of Bayesian inference [5, 6, 4]. A first result concerns the relationship of the phaser to the spiders of CQM:

Theorem. The phaser, when the 1st argument is pure, takes the form of a spider where the ONB in which the spider is expressed arises from diagonalisation of the 2nd argument. Explicitly, setting $|x\rangle = (x_1 \ldots x_{n-1})^T$, we have:

$$\Big(|\psi\rangle\langle\psi|\Big)$$  $$\left(\sum_i x_i^2 |i\rangle\langle i|\right) = |\phi\rangle\langle\phi| \qquad \text{where} \qquad |\phi\rangle :=$$

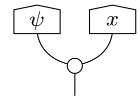

Importantly, these spiders may be expressed in different ONBs which they inherit from the 2nd argument σ, and if we update with nouns which diagonalise in different bases, then the corresponding spiders typically won't commute:

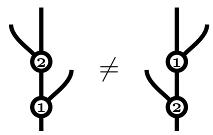

Hence, for the phaser, it is the properties with which the nouns are updated that control commutativity.

The existence of two candidates for meaning updating seems to indicate that there is no unique grey dot. Also, neither the fuzz nor the phaser is a CP-map on the input $\rho \otimes \sigma$. In other words, these update mechanisms are not *internal* to the meaning category. Both of these problems can be overcome by considering the *double density matrices* (DDMs) of [1, 10], which take the form:

$$\sum_{ijk} |\omega_{ik}\rangle|\bar{\omega}_{ik}\rangle|\omega_{jk}\rangle|\bar{\omega}_{jk}\rangle \qquad \text{i.e. diagrammatically}$$

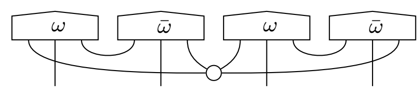

Using spectral decomposition of $\sum_i |\omega_{ik}\rangle\langle\omega_{ik}|$ and map-state-duality this can be rewritten as:

$$\sum_k y_k \left(\sum_i x_{ik} P_{ik} \circ P_{ik}\right) \circ - \circ \left(\sum_j x_{jk} P_{jk} \circ P_{jk}\right) \qquad \text{i.e. diagrammatically}$$

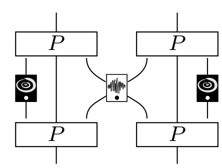

where we indicated indicate the roles of fuzz and phaser. Our general update mechanism including a grey dot merely consisting of wires now is:

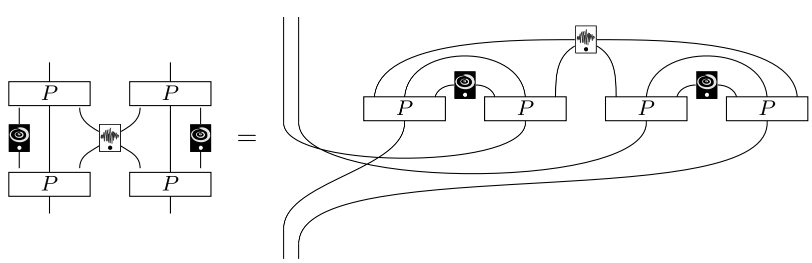

so we can conclude:

Theorem. DMMs enable one to unify fuzz and phaser in a combined update mechanism, where fuzz and phaser correspond to the two modes of mixedness of DDMs. In order to do so, meanings of propositions are generalised to being DMMs, and the update dot is then entirely made up of wires only:

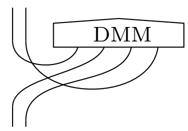

We can still take as meaning category density matrices with CP-maps as processes, since we can indeed think of a DMM as a density matrix of the form:

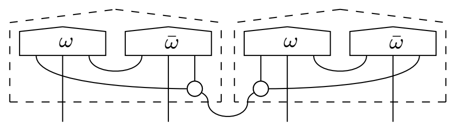

The update dot is a CP-map of the form:

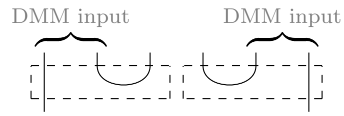

In this way we obtain a meaning category for which updating is entirely internal.

The remainder of the full version `arXiv:1904.03478` provides some examples of this formalism in action, for example, the changing colours of things after over-painting, and how the two modes of mixing, fuzz and phaser, are both needed, for example for describing the meaning of `black metal fan`.

References

[1] D. Ashoush and B. Coecke. Dual Density Operators and Natural Language Meaning. *Electronic Proceedings in Theoretical Computer Science*, 221:1–10, 2016. arXiv:1608.01401.

[2] G. Birkhoff and J. von Neumann. The logic of quantum mechanics. *Annals of Mathematics*, 37:823–843, 1936.

[3] B. Coecke. The mathematics of text structure, 2019. arXiv:1904.03478.

[4] B. Coecke and R. W. Spekkens. Picturing classical and quantum bayesian inference. *Synthese*, 186(3):651–696, 2012.

[5] M. S. Leifer and D. Poulin. Quantum graphical models and belief propagation. *Annals of Physics*, 323(8):1899–1946, 2008.

[6] M. S. Leifer and R. W. Spekkens. Towards a formulation of quantum theory as a causally neutral theory of bayesian inference. *Physical Review A*, 88(5):052130, 2013.

[7] M. Lewis. Compositional hyponymy with positive operators. In *Proceedings of Recent Advances in Natural Language Processing, Varna, Bulgaria*, pages 638–647, 2019.

[8] D. Widdows. Orthogonal negation in vector spaces for modelling word-meanings and document retrieval. In *41st Annual Meeting of the Association for Computational Linguistics*, Japan, 2003.

[9] D. Widdows and S. Peters. Word vectors and quantum logic: Experiments with negation and disjunction. *Mathematics of language*, 8(141-154), 2003.

[10] M. Zwart and B. Coecke. Double Dilation $\neq$ Double Mixing. *Electronic Proceedings in Theoretical Computer Science*, 266:133–146, 2018. arXiv:1704.02309.