Duality of Syntax and Semantics — From the View Point of Brain as a Quantum Computer*

Joseph C.H. Chen and **Manfred Kudlek**
Computer Science Department, University of Hamburg
Vogt-Kölln-Strasse 30,
D-22527 Hamburg, Germany
`joseph@nats.informatik.uni-hamburg.de`
`kudlek@informatik.uni-hamburg.de`

Abstract

In this article, a novel approach to natural language by taking the human brain as a quantum computer is proposed. In this conjecture, symbols are treated as eigenstates with respect to a particular quantum experimental arrangement associated with an *operator* (called a *language formulation*) that corresponds to an *observable* in quantum mechanics. The well-established mathematical formalism of quantum mechanics can be then applied to reasoning as well as natural language processing. The duality of syntax and semantics can be therefore envisaged as an implication of the *Uncertainty Principle.* A miniature bilingual environment is presented to show the feasibility of this idea.

1 Motivation

Taking modern natural sciences into account, any theory of human cognition in general, and natural language ability in particular, can hardly escape from a minimal *physicalist* account — that the physical substrate (the brain) and the functions of human cognition do mesh. In particular, in many theories of mind and language the brain is taken as a passive machine which *obeys* the physical laws (or rules) without exception. An unfortunate consequence of this "objectified" and "neutralized" approach is, of course, that it precludes any *interesting* topic from the so-called scientific discourse, for *interest* (or *intention*) is, by definition, subjective[1]. In this sense, most physicalist theory of mind and language — this includes classical computational approaches — are theories of "zombies."

This is a quite unfortunate situation for natural language processing (NLP). As it happens to be the case, human interest is the centerpiece of human language usage. To wit, any *meaning* in language seems to be rooted in certain intention of the speaker / writer or the listener / reader. Even in seemingly "meaningless" activities (such as babbling or playing with words), intention takes its part, for it is the joy or perplexity of babbling or playing *per se* which may *make sense.*

Indeed, the difficulty of lifeless "theories of zombies" in conventional artificial intelligence (AI) — and connectionism alike — can be largely traced to its root in Cartesian mind-body dualism (Descartes 41), of which the culmination is classical physics. Nevertheless, the advent of quantum mechanics (Dirac 58) has profoundly changed the classical view in physics. In fact, it can be argued that the micro objects at the quantum level are prototypically *mind-like* and *active.*(Stapp 95; Penrose 94) This, in a way, has alleviated the justified pessimism towards mechanistic AI and offered a remedy to passive mechanistic view of brain. We take this view as the starting hypothesis in this article. After all, quantum mechanics is one of the most triumphant theories of physics, it has subsumed chemistry and eventually offered an account bridging fundamental natural phenomena to biology and the human body.

In the following, we will summarize a pro-

---

* Published in *Proceedings of Recent Advances in NLP*, Tzigov Chark, Bulgaria, 5-7 September 2001.

[1] Not only in a mechanistic determinism there is no place for subjectivity. In a stochastic theory, it is as implausible that *interest* or *intention* can be envisaged as a result of passive random processes.

posal of treating language usage as quantum computation. We observe that symbols are invariant throughout a discourse and a symbol does not mesh with other symbols (a symbol σ is *not* not-σ). Therefore symbols are timely *stationary* and *local.* On the other hand, the concept a symbol refers to is in general *dynamic* and *holistic.* It depends on the context — both spatially and temporally. Given this observation, we propose that symbols are to be treated as invariant states (*eigenstates*) of a quantum measurement. The concept (or the state of affairs), on the other hand, is a pre-measurement *superposition* of eigenstates. Knowing symbols are local and concepts are holistic, it suggests that symbols and concept have a sort of *conjugate* relationship — pretty much like the duality of particle (local) and wave (holistic) in quantum mechanics. This will be detailed in the following sections.

2 Cognition according to quantum mechanics

For lack of space we cannot present the entire mathematical and physical backgound. That can be found in textbooks of quantum mechanics, e.g. (Feynman *et al.* 63; Dirac 58).

2.1 State of affairs as a superposition

The idea of quantum computational NLP is to treat a speaker or a listener (in fact, any cognitive agent) as a quantum physical system (Chen 01a). Specifically, there is a symbol inventory (called the *vocabulary*) of the agent, consisting of *eigenstates* of a particular *language formulation* that corresponds to an *observable* S (an *Hermitian* operator, i.e. $S^\dagger \equiv \{S^t\}^* = S$, where ^{t} and * are the transpose and complex conjugate operation in matrix representation, respectively) in the formalism of quantum mechanics. Each eigenstate is associated with a symbol in the inventory. Furthermore, these eigenstates form a complete *orthonormal* basis in a Hilbert space in which the *states of affairs* are to be addressed. Generally speaking, a state of affairs (denoted as $|m\rangle$ — we use Dirac's *bra-ket* notation (Dirac 58) and the Schrödinger picture in the following discussion) is a *superposition* of these basic kets,

$$|m\rangle = \sum_n c_n \, |s_n\rangle \tag{1}$$

where $c_n \in \mathbb{C}$ is a complex number with $c_n = \langle s_n \mid m\rangle$ being the *projection* of $|m\rangle$ on the eigenstate $|s_n\rangle$; $|s_n\rangle$ fulfills the property: $S\,|s_n\rangle = \lambda_n \, |s_n\rangle$ with $\lambda_n \in \mathbb{R}$ being a real number (the *eigenvalue* of $|s_n\rangle$ corresponding to S). For a state of affairs, the probability of finding a particular symbol (by some kind of measurement) corresponding to $|s_i\rangle$ is $P(s_i) = |c_i|^2 / \sum_n |c_n|^2$. For convenience, we can normalize $|m\rangle$ so that it always has the length of unity. Furthermore, we assume that a reasoning based on a state of affairs can be regarded as an undisturbed evolution of the system. This can be described by a *unitary* operator U (a unitary operator is an operator with the property $U^\dagger = U^{-1}$). In a quantum system with constant energy, given an initial state of affairs $|\phi_0\rangle$, the state of affairs at time t is,

$$|\phi(t)\rangle = U \, |\phi_0\rangle = e^{-iHt/\hbar} \, |\phi_0\rangle \,, \tag{2}$$

where H is the *Hamiltonian* operator[2], which is Hermitian; $\hbar$ is the Planck constant divided by 2π. Notice that it is convenient to represent a state of affairs in terms of a *complex-valued vector* (possibly with infinite dimensions). An operator is then a *complex-valued matrix* and the projection $\langle \psi \mid \phi \rangle$ can be treated as the usual inner-product $\langle \psi \mid \phi \rangle \triangleq \sum_n {\psi_n}^* \phi_n$ with ϕ_n (ψ_n) being the n-th component of $|\phi\rangle$ ($|\psi\rangle$) and $*$ being the complex conjugate ($(a+bi)^* = a - bi$ for $a, b \in \mathbb{R}$).

2.2 The uncertainty principle

There is an important implication of our treating a symbol inventory as the eigenba-

---

[2] Hamiltonian is the energy operator. In a closed system, Hamiltonian is a constant matrix.

sis of a particular language formulation operator. In fact, a Hilbert space can be decomposed in different ways. Suppose there are two Hermitian operators S_1 and S_2, each of which can be used as the operator to decompose the space of states of affairs into eigenkets. S_1 and S_2 may not *commute*. That is, there may exist a state of affairs $|\phi\rangle$ such that $S_1 S_2 |\phi\rangle \neq S_2 S_1 |\phi\rangle$, or denoted as the *commutator* $[S_1, S_2] \equiv S_1 S_2 - S_2 S_1 \neq 0$. As in physical systems we assume $[S_1, S_2] = \pm i\hbar I$, where I is the identity operator. Then we have, using straight-forward application of Schwartz's inequality,

$$\Delta S_1 \Delta S_2 \geq \hbar/2. \tag{3}$$

where $\Delta S = \sqrt{\langle (S - \langle S \rangle)^2 \rangle}$; $\langle \cdot \rangle$ is the expectation value. This is the *uncertainty principle* of language formulation. For instance, if S_1 and S_2 are the language formulation operator of English and that of German respectively, it is clear that these two operators do not commute (consider the symbol "`Taube`" in German and "`dove`" or "`pigeon`" in English; or "`fressen`" or "`essen`" in German and "`to eat`" in English).

An important corollary of Equation 3 is an uncertainty relation between *symbols* (in spoken or written language) and *concepts* (as the "real-world referents" of symbols). This can be argued by treating the non-verbal *memory*, whatever it may be, as a patterned representation system embedded in the neuronal quantum mechanical substrate. The memory is therefore a particular sort of language and follows the principle of quantum mechanics.

In particular, if we see syntax as the study of the well-formed formulas of a formal logic or linguistic system — this is coherent with the conventional view of grammar, morphology and phonology — the conventional AI approach to semantic is nevertheless a syntactic approach. Indeed, if one refuses the temptation to "smuggle" interpretation from without (i.e. by "design"), the relationship between concepts is embodied nothing but in an artificial language formalism, in which pure syntactical properties are taken into account. Notice that we deliberately avoid the word *meaning* in discussing semantics. For meaning, we see it as an *active* measurement of a state of affairs. In this regard, the fallacy of conventional computational semantics is that it does not treat meaning as a *meaning giving process*. Instead, it takes meaning as "well defined" "objects" and given relations among "objects." In fact, "well-definedness," *not* probability, is the crucial point in quantum theory vis-a-vis classical deterministic or stochastic theories. According to quantum theory, the properties of a physical system are *not* physically well defined until a measurement is actually performed. This is where intention takes its part[3], which turns *potentialities* into *actualities*.

In this preliminary study, we employ a bilingual miniature environment as a test. In fact, since concepts can be embodied in different (artificial) languages, it can be embodied in a natural language as well. In this way, we have a scenario in which syntax-semantics relationship can be treated as a problem of machine translation.

2.3 Determination of the unitary operator U

In this section, a preliminary algorithm of finding the unitary quantum operator will be presented. The reader should bear in mind, however, that the algorithm presented here is only to show the feasibility of our conjecture. Specifically, it is a *first order* scheme which can be comfortably simulated on a conventional computer and may contain "quick-and-dirty" engineering tweaks. On the other hand, on a fully-fledged quantum computer — the human brain, if our conjecture is correct — there must be some crucial computations which are beyond the scope of classical computation (Steane 98). The reader is therefore urged *not* to judge our proposal

---

[3] In fact, intention is as well a process, not an object.

only from the limitedness of the algorithm but to focus on the underlying interpretation despite of the superficial similarity of our approach and classical connectionism or stochastic frameworks. That said, let us present our miniature NLP problem.

In a normal discourse we use *orthographic* utterances, that is, strings consisting of symbols in the vocabulary inventory to convey ideas. An orthographic string has to be converted to a quantum state of affairs complying with Equation 1 (i.e a complex-valued vector). This can be a very complicated procedure. In the implementation discussed in this article, however, a simplified model is presented in which an initial state of affairs (associated with a sentence) is represented by,

$$|\phi_i\rangle = \sum_{k=1}^{m} e^{i(k-1)\theta_0}\,|s_k\rangle, \qquad (4)$$

where $\theta_0 = 2\pi/(m+2)$ with m being the length of the sentence; each of the symbols in the sentence is represented by its corresponding eigenstate. Specifically, the word order is encoded in the *argument* or the *phase* of the complex components (remember that any complex number z can be written in the form $z = re^{i\theta}$ with $r \geq 0, \theta \in \mathbb{R}$). The goal is then to find a suitable Hamiltonian so that a particular reasoning can transform the initial state of affairs into the desired end state of affairs. Furthermore, we assume that the result is to be taken at an arbitrary time point. In this case, $\hbar$ and t can be absorbed. That is, $U = e^{-iH'} = e^{-iHt/\hbar}$. For convenience and without losing generality, we will discuss H' in place of Hamiltonian and denote it as H hereafter.

Since H is an Hermitian (i.e. $H^\dagger = H$) matrix, there are in total n^2 free *real* parameters to be determined, provided the size of the vocabulary is n. The optimal parameters can be found using a standard optimization algorithm. Specifically, we use the conjugate gradient method (Press *et al.* 92) with the cost function defined as,

$$err(H) = \sum_{(\phi_t,\phi_i)\in T} \left| \left\langle \phi_t^k \,\middle|\, \phi_o^k \right\rangle \right|^2, \qquad (5)$$

where T is a set of training pairs (ϕ_t, ϕ_i); $|\phi_t\rangle$ and $|\phi_i\rangle$ are the target and input state of affairs respectively. Moreover, the output state $|\phi_o\rangle$ is related to $|\phi_i\rangle$ as follows

$$|\phi_o\rangle = U|\phi_i\rangle = e^{-iH}|\phi_i\rangle, \qquad (6)$$

and $|\phi_t\rangle$ is also prepared from the orthographic string according to Equation 4.

In the decoding phase, the end state of affairs has to be converted to its orthographic form. In general, this should be done by subject the end state to another unitary operation and observing the state in successive time intervals. In this initial study, however, a heuristic combinatorial algorithm is used to construct the orthographic sentence instead. Specifically, every candidate sentence ψ is given a score ($score(\psi) = |\langle\psi \mid \varphi\rangle|$), where φ is the end state. In the ideal case, the inner product should be unity (1) for a perfect candidate. The search for the best candidate is done according to the following algorithm,

```
0. Normalize the end state of affairs;
   Theta:=0.01;
1. Build a set S of all symbols with
   absolute value >= Theta;
2. Calculate the score of each permutation
   in S; notice the one with best score;
3. Theta := Theta + 0.01;
4. If Theta <= 0.4 goto step 1;
5. Output the permutation with best score.
```

3 Preliminary experiments on machine translation

In this section, the constructed unitary operator U is applied on the machine translation of a bilingual corpus. The experimental data are based on the corpus proposed by Chalmers (Chalmers 90). We have redone his initial experiment with our framework and achieved both accuracy and recall of 100% (compared with 65% generalization accuracy reported in Chalmers' paper). For a more realistic experiment, we used the vocabulary summarized in Table 1 and generated an extended bilingual corpus. Specifically, both the active and passive utterances are translated to German resulting in 480 English-German bilingual sentence pairs in total. All the verbs are correctly conjugated.

There are in total 19 (20) symbols in the English (German) vocabularies. The corpus is not trivial. For instance, there is a separable German verb (*umbringen*[4] — to kill) in the corpus which may be mapped to two eigenstates in some formulations.

Vocabulary Category	English Instances	German Instances
Person	john michael helen diane	john michael helen diane
Verb	kill love * betray hit	umbringen† lieben verraten* schlagen
Verb conj.	kills loves betrays hits	bringt† liebt verraetet schlaegt
Past Participle	killed loved betrayed hit*	umgebracht geliebt verraten* geschlagen
Conjunction	and	und
Misc.	is are by	wird werden von um†

Table 1: Vocabulary used in the bilingual corpus. Words marked with * are "homonyms" which are represented by identical eigenstate in the vocabulary. † German verb *umbringen* is a separable verb.

Seventy-eight sentence-pairs are chosen randomly as training set (16% of the corpus) and the remaining 408 sentence-pairs are reserved as test set. We then used conjugate gradient method (Equation 5 as the cost function) to determine the unitary operator U based on the training set. The total number of free parameters is therefore $39^2 = 1521$. Once the unitary operator is obtained, each utterance in the test set is constructed according to Equation 4 and subject to U, resulting in a corresponding end state-of-affairs. The end state-of-affairs is decoded using the combinatorial algorithm described in Section 2.3 so that an orthographic sentence can be reconstructed accordingly. In a typical experiment, the correctness on the training set is 93.6% and the generalization rate on the test is 88.8%. If the correctness of words instead of that of sentences is counted, the correctness of the training set rises to 97.8% and the generalization accuracy of the test set is 95.8%. Given the small size of the training set, the generalization rate is very impressive. (In the percentages we omitted error intervals.)

Nevertheless, there are some sentence-pairs which are not correctly learned either in the training or in the test set. For example, in the test set, the pair

```
helen is killed by michael and diane ->
helen wird von michael und umgebracht* diane*
(helen wird von michael und diane umgebracht)
```

which is not correctly learned, is shown in Figure 1. Incorrectly decoded words are marked with * (this includes incorrect word order as well). The target is shown in the parentheses. As can be seen in the figure, the output end state of affairs is largely similar to that of the target. The error is mostly due to phase-shifts, resulting in an incorrect word order. If the order is swapped, the errors are corrected. In fact, if the decoded sequences are permuted only once (by swapping the positions of exactly two of the symbols), we achieve an accuracy of 100% on the training set. On the test set, the accuracy is 98.5%.

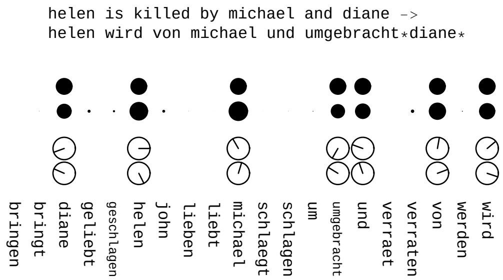

Figure 1: An example of the test set in the bilingual corpus which is not correctly decoded. The first two rows are the absolute squares (in terms of area) of the target and the output state respectively; the last two rows are the phase of the target and the output respectively. Remember that any complex component c_n can be written as $r_n e^{i\theta}$ with $r_n \geq 0$ and $0 \leq \theta < 2\pi$.

A glance at the remaining errors (6 sen-

---

[4] The prefix of a German separable verb is placed at the end of the sentence and represented by two independent symbols in conjugated form. For example, `diane bringt john um <-> diane kills john`.

tences), we notice that they are all of the form as shown in the following:

```
diane kills helen ->
diane bringt john* helen* um-
(diane bringt helen um)
```

of which the end state of affair is shown in Figure 2. The error of this example is due to unwanted residues of eigenstates. This kind of error can be removed by raising the threshold in the combinatorial decoding process. For example, if the threshold is set to 0.1, all the error of this kind can be alleviated.

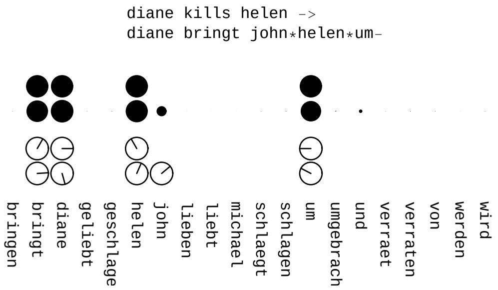

Figure 2: An example of the test set in the bilingual corpus which is not correctly decoded. The error is mainly due to residue of irrelevant eigenstates.

This example also shows an interesting "bias" of the system to convict `john` as killer. In fact, this comes with no surprise if we take a closer look at the training set. In the training set there are 20 sentences which are about "killing." In these scenarios, `john` kills 9 times and is killed 6 times, which is the most frequent killer except `diane` (`michael` kills 7 times and is killed 9 times. `helen` kills 7 times and is killed 8 times. `diane` kills 7 times and is killed 3 times.) The poor `john` seems to have become the natural "black sheep" owing to the unbalanced training set.

An interesting but non-trivial "by-product" of this experiment is that one can use the architecture to compile a bilingual dictionary of the miniature languages. This can be done by using the lexical list of English as input and looking at the end state of affairs in German. The result is shown in Figure 3. As can be seen in the map, the English words are largely associated with their German translations. Interestingly, personal names and auxiliary words (`is`, `are`, `by`) are mapped distributively to several German words. The German counterparts of the English personal names are nevertheless the most activated. The auxiliary words, on the other hand, show a sort of "template" relationship, which is basically a *many-to-many* mapping. These are desirable results that correspond well to our intuitive understanding of language usage and "meaning" of words.

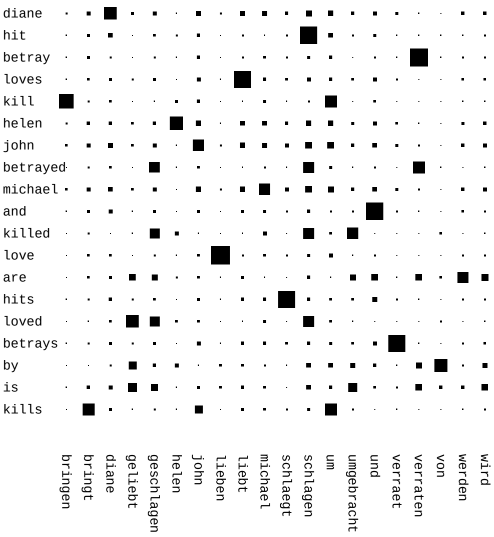

Figure 3: A English-German dictionary map. The absolute square of each component is represented by the area of the little square

4 Discussion and conclusion

The treatment of cognition and language as a quantum system is motivated more or less by the philosophical discontent with the classical mechanistic view of AI (this includes connectionism) and its inability to offer a coherent view of physical substrate and mental activities (other than blatant dualism) (Chen 01b). In this regard, our proposal is a radical departure from the conventional AI and linguistic view of symbols, which, to wit, have to be implemented as well-defined physical objects (Newell & Simon 76). It is at the same time a departure from connectionism and stochastic approaches, which treat symbols (or representation) as nothing but continuous real parameters. Unlike classical-physics-based physicalism, in quantum computation there is no need of mysterious postulate of "emergence" of active mind out of passive, mechanistic matter (Churchland & Sejowski 92) to account for *discreteness* and *activeness* aspects of language and logic. In contrast, the activeness of meaning and intention is embedded in the active quantum-measurement. Nevertheless, our approach remains *physicalist* and *naturalist*.

In fact, there are recently proposals that quantum effects should be taken into account in explaining consciousness (Hameroff 98). This agrees with our proposal in that consciousness or intention is embedded in the active measurement of a quantum system. Moreover, if our conjecture is correct, intention (if treated as an objectified entity) must subject to the uncertainty principle as well. This render intention *per se* — as it truly is a holistic process and *any* measurement yields only localized results — beyond the access of our language in particular and representation in general.

Last but not least, the preliminary experiment shows that a quantum mechanical architecture can accomplish miniature NLP tasks quite successfully. This suggest that a quantum mechanical framework may be applied to practical problems without waiting for a true quantum computer. For instance, one may build a hybrid model in which classical symbolic computation and/or statistical-connectionist modules work together with a quantum mechanical "arbitration" module. The latter is in charge of crucial decisions. This is similar to a scenario in which a human is assisted by computer programs or other knowledge sources to make decisions. In fact, if the conjecture in this paper is justified, the human brain may very likely work in a similar classical-quantum hybrid mode, since classical mechanics is a limiting case of quantum mechanics.

References

(Chalmers 90) D.J. Chalmers. Syntactic transformations on distributed representations. *Connection Science*, 2(1&2):53–62, 1990.

(Chen 01a) Joseph Chen. A quantum mechanical approach to cognition and representation. In *Toward a Science of Consciousness 2001*, Skövde, Sweden, 2001.

(Chen 01b) Joseph C.H. Chen. *Quantum Computation and Natural Language Processing.* University of Hamburg, Germany, 2001. Ph.D. thesis in preparation.

(Churchland & Sejowski 92) P. S. Churchland and T. J. Sejowski. *The Computational Brain.* MIT Press, Cambridge, MA, 1992.

(Descartes 41) René Descartes. *Meditations on first philosophy.* 1986 Cambridge University Press, Cambridge [Cambridgeshire] ; New York, 1641.

(Dirac 58) P. A. M. Dirac. *The principles of quantum mechanics.* International series of monographs on physics. Clarendon Press, Oxford,, 4th edition, 1958.

(Feynman *et al.* 63) Richard Phillips Feynman, Robert B. Leighton, and Matthew L. Sands. *The Feynman lectures on physics.* Addison-Wesley Pub. Co., Reading, Mass., 1963.

(Hameroff 98) Stuart Hameroff. Quantum computation in brain microtubules? the Penrose-Hameroff "Orch OR" model of consciousness. *Philosophical Transactions Royal Society London (A)*, (356):1869–1896, 1998.

(Newell & Simon 76) Allen Newell and Herbert A. Simon. Computer science as empirical inquiry: symbols and search. *Communications of the ACM*, 19(3):113–126, 1976.

(Penrose 94) Roger Penrose. *Shadows of the mind.* Oxford University Press., New York, Oxford, 1994.

(Press *et al.* 92) William H. Press, S.A. Teukolsky, W.T. Vetterling, and B.P. Fl annery. *Numerical Recipes in C.* Cambridge University Press, 1992.

(Stapp 95) Henry Stapp. Why classical mechanics cannot naturally accommodate consciousness but quantum mechanics can. *PSYCHE*, 2(5), 1995. http://psyche.cs.monash.edu.au/v2/psyche-2-05-stapp.html.

(Steane 98) Andrew Steane. Quantum computing. *Reports on Progress in Physics*, 61(5):117–173, 1998.