Quantum Theory:
From the whole to the parts,
magic squares & shadows of infinity

The relation between quantum physics (events and their frequencies, to the left) and quantum theory (to the right) can be imagined as in this artwork by Kumi Yamashita. This Habilitation Thesis sheds light on some aspects of quantum theory: on going from the whole to parts, on quantum magic squares, and on shadows of infinity.

Professorial Dissertation to obtain the Venia Docendi
at the Faculty for Computer Science, Mathematics and Physics
at the University of Innsbruck
in the Habilitation subject "Theoretische Physik - Theoretical Physics"
presented by Dr. Gemma De las Cuevas
Innsbruck, February 2022

Habilitationscrift zur Erlangung der Lehrbefugnis
für das Fach "Theoretische Physik - Theoretical Physics"
an der Fakultät für Mathematik, Informatik und Physik
der Leopold-Franzens-Universität Innsbruck
vorgelegt von Dr. Gemma De las Cuevas
Innsbruck im Februar 2022

Prologue

This Habilitation concerns quantum theory, and I like to imagine its relation to quantum physics with the artwork by Kumi Yamashita of Fig. 1. On the left we have quantum physics, understood as experimental quantum physics, which consists of events, potentially repeated, giving rise to relative frequencies of these events. Which events are observed is of course all but accidental—it is guided by quantum theory—but ultimately they are events. On the right we have the theory of quantum physics, whose relation to the physical world is very mysterious to me. That relation between the abstract and actual is particularly thorny in the case of quantum theory, given the disagreement on the metaphysics of the theory—does quantum theory tell us that we live in a multiverse, and that there are virtually infinitely many copies of myself? Or do we "just" live in one non-local-realistic universe? And: how is it possible that we disagree on such transcendental matters? While very important, these questions will be set aside in this Habilitation.

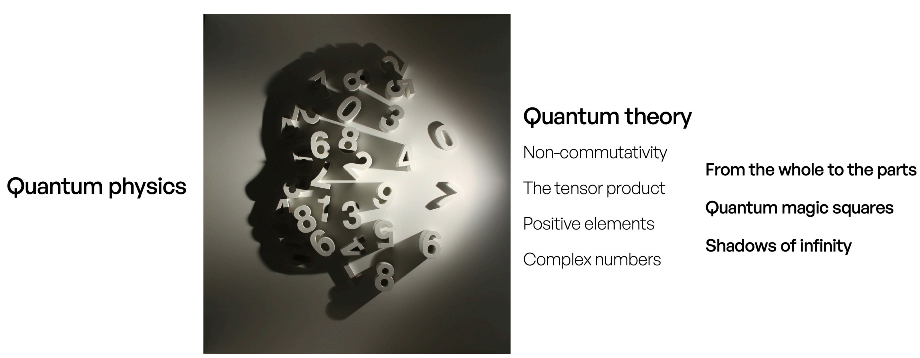

Fig. 1: Quantum theory (the abstract world) can be imagined to shed a very special light on the physical world, namely give rise to surprising events. This Habilitation investigates the interplay of four aspects of quantum theory—non-commutative spaces, the tensor product, the role of positive elements, and the role of complex numbers—through "From the whole to the parts", "Quantum magic squares", and "Shadows of infinity".

Instead, I would like to focus on some mathematical aspects of quantum theory, of which I want to highlight four. The first is the fact that quantum systems are described with *non-commutative* structures. Discovering non-commutativity can be imagined as discovering a new universe (Fig. 2): Classical physics corresponds to level 1 of this universe, whereas quantum physics corresponds to any level higher than one. The level is to be understood as the internal dimension of the system. For example, a qubit is a system at level 2, because it is represented by a 2×2 matrix, whereas a qudit is a system at level d, because it is represented by a $d \times d$ matrix. A classical system is represented by a 1×1 matrix, i.e. by a number—this would describe a random classical variable. The need to represent the state of the system by a matrix testifies to the non-commutative nature of the theory that we are trying to highlight.

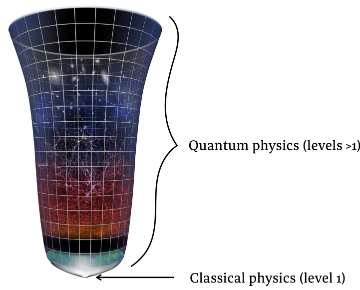

Fig. 2: Discovering quantum physics is like discovering a new universe.

The second aspect is the tensor product, which gives rise to a new compositional universe (Fig. 3). Describing how to compose is as fundamental as describing single systems, and I believe it is a misconception of reductionism to overestimate the importance of single systems. In the case of quantum theory, much of the 'magic' and 'mystery' come from composing with the tensor product instead of the Cartesian product. Namely, the joint system of two quantum systems (two 'universes', in our running metaphor) is described by their tensor product. Physicists tend to imagine the tensor product so that if we pick an orthonormal basis for the first universe, for every element of this basis there is an entire second universe attached to it. A more formal definition of the tensor product is fairly complicated, and invokes the universal property of the tensor product in category theory.

The third aspect is that part of this universe is positive, and these positive elements play a distinguished role (Fig. 4). The positive elements form a cone—namely the cone of positive semidefinite matrices—which is much harder to describe than the underlying 'universe', which forms a vector space. In fact, the state of a quantum system is described by a positive semidefinite matrix of trace 1, so for every level of this universe we do not have a cone but a convex set. Now, these cones

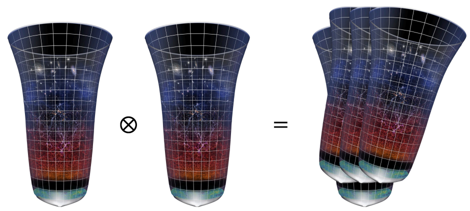

Fig. 3: The tensor product gives rise to a new compositional universe.

interact in a very rich way with the compositional structure of the universes: the tensor product of the cones is not the cone of the tensor product (because the tensor product interacts nicely with vector spaces, not with cones), and the convex combination of cones is not the cone of composite system system either (i.e., there are entangled states). In this Habilitation we will explore this very fertile interaction from various angles.

Fig. 4: The cone of positive elements play a distinguished role in quantum theory, because they are used to describe quantum states.

The fourth aspect is the fact that this universe uses complex numbers (Fig. 5): It has a real and an imaginary part, and so do its positive elements. The composition mixes up the real and imaginary part in a 'consistent' way as far as the number of parameters is concerned: A $d \times d$ complex positive semidefinite matrix has d^2 free real parameters, and the composition of a quantum system at level d and one at level s is described by a $ds \times ds$ matrix, which has $(ds)^2 = d^2s^2$ real parameters. In this Habilitation we will challenge this aspect by considering a *hyperreal* and a *hyperimaginary* part, and we will show that some long-standing problems in quantum theory can be solved over the hypercomplex numbers.

More generally, in this Habilitation we investigate the interplay of these four ingredients from three perspectives: First we will go from the whole to the parts, from many different perspectives (Section 1), then we will study magic squares and their quantum cousins (Section 2), and finally we will investigate shadows of infinity in the problem of tensor stable positivity (Section 3).

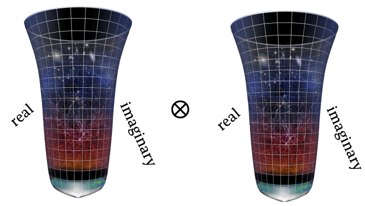

Fig. 5: This universe is complex—it has a real and an imaginary part—, and this interacts nicely with the composition structure.

A faithful companion through our journey will be *free semialgebraic geometry*, which often studies the same mathematical objects as quantum theory but from a different perspective (Fig. 6). 'Free' means noncommutative because it is *free of the commutation relation*, that is, it studies noncommutative versions of semialgebraic sets. Since quantum information theory is a noncommutative generalisation of classical information theory, 'free' is naturally linked to 'quantum'. Moreover, positivity (and convexity) play a very important role in both fields, since semialgebraic geometry examines questions arising from nonnegativity, such as polynomial inequalities, whereas positive elements play a distinguished role in quantum theory, as we highlighted above.

While the two disciplines often study the same objects, they tend to do so from different angles, and hence ask different questions. Generally speaking, free semialgebraic geometry tends to study the sets (their geometrical structure, characterisation, etc), whereas quantum theory puts the emphasis on the *elements* of these sets (because they correspond to quantum states). For example, in quantum theory, given an element of a tensor product space one wants to know whether it is positive semidefinite or entangled, and how this can be efficiently represented and manipulated, whereas in free semialgebraic geometry one characterises the geometry of the set of all such elements. I have had the pleasure to collaborate with my colleague Tim Netzer in recent years—an expert on free semialgebraic geometry—, and of exploring the richness of combining these two perspectives. This connection to free semialgebraic geometry will be a thread seaming the various themes of this Habilitation, and we have written an invitation to the intersection of both topics in P13.

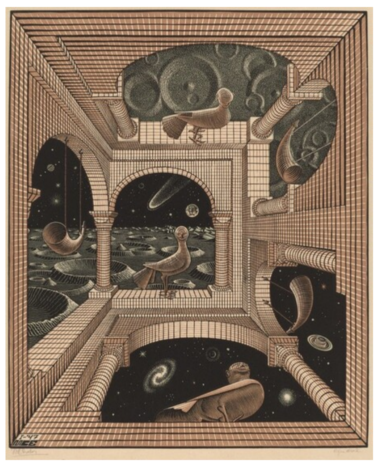

Fig. 6: Free semialgebraic geometry and quantum information theory often look at similar landscapes from different perspectives, as surveyed in P13 and depicted in the fantastic world by M. C. Escher.

List of works

This Habilitation is based on the following works, which are explained in Chapter 1 and included in Chapter 2:

P1. Tensor decompositions on simplicial complexes with invariance — Page 34
G. De las Cuevas, M. Hoogsteder Riera and T. Netzer[1]
Submitted to J. Symb. Comp. arXiv:1909.01737 (2019)

P2. Approximate tensor decompositions: disappearance of many separations — Page 67
G. De las Cuevas, A. Klingler and T. Netzer[1]
J. Math. Phys. **62**,093502 (2021) (Editor's pick). arXiv:2004.10219

P3. Polynomial decompositions with invariance and positivity inspired by tensors — Page 93
Gemma De las Cuevas, Andreas Klingler, Tim Netzer[1]
Submitted to Comm. Math. Phys. arXiv:2109.06680 (2021)

P4. Separability for mixed states with operator Schmidt rank two — Page 133
G. De las Cuevas, T. Drescher, T. Netzer[1]
Quantum **3**, 203 (2019). arXiv:1903.05373

P5. Mixed states in one spatial dimension: decompositions and correspondence with nonnegative matrices — Page 148
G. De las Cuevas and T. Netzer[1]
J. Math. Phys. **61**, 041901 (2020). arXiv:1907.03664

P6. Optimal bounds on the positivity of a matrix from a few moments — Page 175
G. De las Cuevas, T. Fritz and T. Netzer[1]
Comm. Math. Phys. **375**, 105 (2020). arXiv:1808.09462

P7. Irreducible forms for Matrix Product States: Theory and Applications — Page 198
G. De las Cuevas, J. I. Cirac, N. Schuch, D. Perez-Garcia
J. Math. Phys. **58**, 121901 (2017). arXiv:1708.00029

P8. Continuum limits of Matrix Product States — Page 213
G. De las Cuevas, N. Schuch, D. Perez-Garcia, J. I. Cirac
Phys. Rev. B **98**, 174303 (2018). arXiv:1708.00880

P9. Generalised ansatz for continuous Matrix Product States — Page 221
M. Balanzó-Juandó and G. De las Cuevas
Phys. Rev. A. **101**, 052312 (2020) (Editor's suggestion). arXiv:1908.09761

P10. Cats climb entails mammals move: preserving hyponymy in compositional distributional semantics — Page 233
G. De las Cuevas, A. Klingler, M. Lewis and T. Netzer[1]
Accepted as a talk for SemSpace2020 and submitted to J. Cog. Sci.
arXiv:2005.14134 (2020)

P11. Quantum magic squares: dilations and their limitations — Page 276
G. De las Cuevas, T. Drescher and T. Netzer[1]
J. Math. Phys. **61**, 111704 (2020) (Featured by the Editor). arXiv:1912.07332
See AIP Scilight, UIBK news, Tendencias21, El Periodico.

P12. Halos and undecidability of tensor stable positive maps — Page 293
Mirte van der Eyden, Tim Netzer, Gemma De las Cuevas
Submitted to a Special Issue on 'Foundational Structures in Quantum Theory' (Guest Editors: Giulio Chiribella, Bob Coecke, Teiko Heinosaari, Ana Belén Sainz and Robert Spekkens) of J. Phys. A. arXiv:2110.02113 (2021)

P13. Quantum Information Theory and Free Semialgebraic Geometry: One Wonderland Through Two Looking Glasses — Page 324
G. De las Cuevas and T. Netzer[1]
IMN Nr. 246 (2021). arXiv:2102.04240

[1]Alphabetical order.

Contents

Prologue **i**

List of works **vi**

1 Shedding some light on Quantum Theory **1**

1 From the whole to the parts . . . 1

1.1 How do the parts form the whole? . . . 4

1.2 If the whole is invariant, are the parts invariant too? . . . 7

1.3 If the whole is positive, are the parts positive too? . . . 8

1.4 And in the approximate case? . . . 14

1.5 What about other worlds with the same parts–whole relation? 16

1.6 Given some shadows of the whole, is the whole positive? . . . 18

1.7 Which wholes can be divided into arbitrarily many parts? . . . 21

1.8 Which composition rules preserve the order of the parts? . . . 24

2 Quantum magic squares . . . 25

3 Shadows of infinity . . . 27

Bibliography . . . 30

2 Publications & Preprints **34**

P1 Tensor decompositions on simplicial complexes with invariance . . . 34

P2 Approximate tensor decompositions . . . 67

P3 Polynomial decompositions with invariance and positivity . . . 93

P4 Separability for mixed states with operator Schmidt rank two . . . 133

P5 Mixed states in one spatial dimension . . . 148

P6 Optimal bounds on the positivity of a matrix from a few moments . . . 175

P7 Irreducible forms for Matrix Product States: Theory and Applications . 198

P8 Continuum limits of Matrix Product States . . . 213

P9 Generalized ansatz for continuous Matrix Product States . . . 221

P10 Cats climb entails mammals move . . . 233

P11 Quantum magic squares 276
P12 Halos and undecidability of tensor stable positive maps 293
P13 Quantum information theory and free semialgebraic geometry 324

Thank you **353**

Chapter 1

Shedding some light on Quantum Theory

This Habilitation investigates the interplay of four aspects of quantum theory (non-commutativity, the tensor product, positive elements and complex numbers) from three perspectives: From the whole to the parts (Section 1, based on P1, P2, P3, P4, P5, P6, P7, P8, P9 and P10), quantum magic squares (Section 2, based on P11), and shadows of infinity (Section 3, based on P12). In addition, P13 provides an overview of some of these results in the form of an invitation to the topic.

1 From the whole to the parts

The relations between the whole and its parts are a 'classical' topic in metaphysics, and, when constraied to the mathematical realm, concern the study of compositionality —how objects compose and decompose. Composing is usually the 'easy' direction; it is the constructive direction. The inverse problem is decomposing, which could be concerned, e.g., with which properties of the whole can be transferred to the parts. Two notes: First, in the abstract world, the only object whose whole can be put in one-to-one correspondence with a proper part is infinity—this is in fact a defining property of infinity. In our work we will only be concerned with finite systems. The second note concerns the unique and very surprising relations between the parts and the whole in quantum theory: in a maximally entangled state, we have perfect knowledge of the whole and maximal lack of knowledge of the parts. In our work, we investigate the relation between the parts and the whole in a general way for vector spaces composed with the tensor product, with positivity cones playing a distinguished role. This setting includes the description of quantum many-body systems as a special case, but also that of multivariate polynomials, as we will see. We will be guided by the following questions:

Q1. *How do the parts form the whole?* (Section 1.1)

For a given whole, we will describe any decomposition using a weighted simplicial complex Ω. This includes the tensor rank decomposition and the linear decomposition, but many more too.

Q2. *If the whole is invariant, are the parts invariant too?* (Section 1.2)

This initiates the study of which properties of the whole can be transferred to the parts. To start with, we consider invariance under the permutation of some parts, and ask whether there is an expression of the parts that certifies this invariance. We give a sufficient condition for the existence of this certificate.

Q3. *If the whole is positive, are the parts positive too?* (Section 1.3)

For several notions of positivity (i.e. several cones), we ask whether the parts can contain a certificate that the whole is in a cone. The answer is, generally, 'yes' but at a very high price—where 'price' means rank, i.e. number of terms in the decomposition. We will consider invariant and positive wholes, i.e. this question combined with Q2. P1 contains the framework that addresses Q1, Q2 and Q3.

Q4. *What if the parts only approximate the whole?* (Section 1.4)

Does the answer to Q2 and Q3 change if the parts give rise to a whole which is ε-close to the original one? Yes, it does: certifying positivity in the parts can be much easier in the approximate case. This question is addressed in P2, leveraging P1's framework.

Q5. *What about other worlds with the same parts-whole relation, such as multivariate polynomials?* (Section 1.5)

Since multivariate polynomials are also described by the tensor product of vector spaces, we apply our entire framework—which is inspired by tensor decompositions—to polynomials, and transfer many results. This question is addressed in P3 (leveraging P1's framework and P2's results).

Q6. *Given some shadows of the whole, is the whole positive?* (Section 1.6)

By 'shadows' we mean a few moments of a matrix. So the question is: Given a few moments $\mathrm{tr}(M^k)$ for $k = 1, \ldots, m$ of a matrix M, can we tell at most

Fig. 1.1: P1 provides a new framework to decompose objects of tensor product spaces in terms of their parts, including invariance and positivity. P2 addresses the approximate case, and P3 applies the framework to polynomials. This *Gala of spheres*, by Dalí, is a remarkable composition of a whole (Gala's bust), or decomposition into some surprising parts.

and at least how far it is from the cone of positive semidefinite matrices? We address this question in P6.

Q7. *Which wholes can be divided into arbitrarily many parts?* (Section 1.7)

Namely, which states have a continuum limit? We address this question for Matrix Product States in P8, which requires a generalisation of the canonical form and its fundamental theorem provided in P7. Since our results imply that not every continuum limit of a Matrix Product State can be expressed as a continuous Matrix Product State, we propose a generalisation of the latter in P9.

Q8. *Which composition rules preserve the order of the parts?* (Section 1.8)

If the parts are in a cone, they build a partial order, but this relation is generally not preserved by the composition of the parts. In P10 we propose a composition rule that preserves this relation. This is relevant for recent approaches to (quantum) natural language processing, which represent the meaning of words by positive semidefinite matrices, and hyponymy by the partial order of positive semidefinite matrices. Our composition rule, thus, preserves hyponymy.

1.1 How do the parts form the whole?

If the parts are vector spaces $V_1, \ldots V_n$, giving rise to the composite vector space $V = V_1 \otimes \ldots \otimes V_n$, and the whole is an element $v \in V$, the question is:

How can v be decomposed into its parts, that is, into vectors of V_i?

Clearly, v can always be decomposed as

$$v = \sum_{\alpha=1}^{r} a_\alpha \otimes b_\alpha \otimes \ldots \otimes z_\alpha, \tag{1.1}$$

but bear in mind that this decomposition is highly non-unique. What is unique is the minimal such r, called the *tensor rank* of v. If we only had three parts ($n = 3$), we could represent this decomposition as a full simplex,

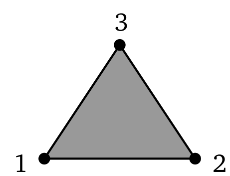

where the filled face indicates that the summation index α is shared by the 3 parts.

But there are many other ways to decompose v. The parts could be connected 'in a line',

$$v = \sum_{\alpha_1,\alpha_2,\ldots,\alpha_{n-1}=1}^{r} a_{\alpha_1} \otimes b_{\alpha_1,\alpha_2} \otimes c_{\alpha_2,\alpha_3} \otimes \ldots \otimes z_{\alpha_{n-1}} \tag{1.2}$$

represented as

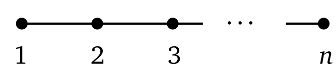

While this is the natural decomposition to describe quantum many-body systems in one spatial dimension (with open boundary conditions), it is essentially absent in mathematics. The minimal number of terms for the 'line decomposition' is called the operator Schmidt rank, and it can be much smaller than the tensor rank.

Or the parts of v could be connected 'in a circle',

$$v = \sum_{\alpha_1,\alpha_2,\ldots,\alpha_n=1}^{r} a_{\alpha_n,\alpha_1} \otimes b_{\alpha_1,\alpha_2} \otimes \ldots \otimes z_{\alpha_{n-1},\alpha_n}$$

represented as

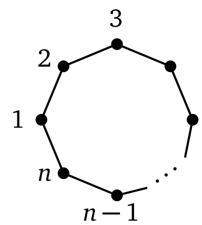

This decomposition would be relevant for the description of one-dimensional quantum many-body systems with periodic boundary conditions, and it is unclear what its minimal number of terms is called, and how it relates to the ranks of the other decompositions.

Or v could be decomposed as

$$v = \sum_{\alpha,\beta,\gamma=1}^{r} a_\alpha \otimes b_{\alpha,\beta} \otimes c_\beta \otimes d_{\beta,\gamma} \otimes e_\gamma \tag{1.3}$$

represented as

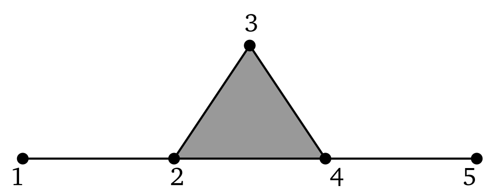

Again it is unclear what the minimal number of terms r is called, and how it relates to the other examples.

Or, if v were bipartite ($n = 2$), its parts could be connected by a double edge,

$$v = \sum_{\alpha,\beta=1}^{r} a_{\alpha,\beta} \otimes b_{\alpha,\beta} \tag{1.4}$$

represented as

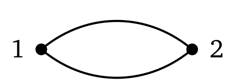

Such decompositions are important for positive decompositions with invariance (Q3), as we will see later.

After these preparatory examples, we can address Q1: How do the parts form the whole? P1 provides a framework to put all such decompositions, and more, under one umbrella given by a *weighted simplicial complex* Ω. The central idea is that *the parts, i.e. individual vector spaces, be associated to the vertices of* Ω*, and the*

the summation indices be associated to the maximal facets of Ω— we have been using this framework in the examples and representations above. A weighted simplicial complex is a "well-behaved" weighted hypergraph. That is, given a set of vertices V, the facets F are elements of its power set $\wp(V)$—for example, the simplex of page 4 is described by a facet with vertices $\{1,2,3\}$. Since the simplicial complex is weighted, these facets can have multiplicity—for example, in the double edge of page 5, the facet $\{1,2\}$ has multiplicity 2.

Moreover, for every simplicial complex Ω, the minimal number of terms in the Ω decomposition defines the rank_Ω. The tensor rank decomposition corresponds to the case when Ω has a full simplex, i.e. a facet containing all vertices, the circle decomposition to the case when Ω is the circle graph, the decomposition of (1.3) to Ω being the hypergraph of page 5, and the double edge decomposition to the case when Ω consists of two vertices and a double edge.

What can we do with this framework? We can put the knowledge coming from mathematics (mainly concerning the tensor rank or its symmetric version) and that of quantum many-body systems (mainly concerning the operator Schmidt rank or its translationally invariant version) under one umbrella, and compare ranks and transfer results (as done in P1). More generally, we can study questions such as Q2 and Q3, as we will do below. But let us first consider the special case of two parts.

The case of two parts

One final remark: If v only has two parts, this framework is unnecessary (apart from the double edge case, i.e. the multiplicity of the edge, which will be relevant for Q3). In the bipartite case, the minimal number of terms—the rank—fully characterises the dependence of the whole on its parts, in the sense that v is a sterile concatenation of its parts if and only if the rank is 1.[1] The rank is easy to compute, and it is impossible to exaggerate its importance across the natural sciences and mathematics. But the simplicity of the bipartite case is misleading, and it may be an accident of the number two. For three parts it is no longer true that v is a sterile concatenation of all of its parts if and only if the tensor rank is 1. Other examples where switching from 2 to 3 entails a jump in complexity include the famous 2SAT versus 3SAT problems—the former can be solved in polynomial time, whereas the latter is NP-complete.[2]

[1] The rank is a very non-smooth measure of this fact, as it can grow arbitrarily by letting the parts interacts a tiny bit—an ε. This may partly explain the answer to Q4.

[2] NP is the class of decision problems that can be solved in polynomial time by a non-deterministic Turing machine. NP-complete means that the problem is in NP, and that it admits a polynomial-time reduction from any other problem in NP.

1.2 If the whole is invariant, are the parts invariant too?

First of all, what does it mean for the whole to be invariant? We say that v is G invariant (where G is a group) if $gv = v$ for all elements g in G, and gv is the permutation of some parts in v. If G is full permutation group S_n, to say that v is G invariant is to say that v be fully symmetric or permutationally invariant. On the other hand, if G is the cyclic group C_n, to say that v is G invariant is to say that v be translationally invariant. Note that this is an "external symmetry" of v, as it permutes its parts, and does not refer to the internal symmetries of the elements of v, e.g. that they be $SU(2)$ or $\mathbb{Z}_2$ symmetric. These external symmetries are the only symmetries we will study here.

Now, what does it mean for the parts to be invariant? It means that that the parts contain a 'certificate' of this invariance, i.e. that each local tensor is G invariant. For example, the decomposition

$$v = \sum_{\alpha=1}^{r} a_\alpha \otimes a_\alpha \otimes \ldots \otimes a_\alpha \tag{1.5}$$

is S_n invariant—if v admits this decomposition, then v is S_n invariant, because any permutation of the above expression leads to itself. (The minimal number such r is called the symmetric tensor rank). Or the decomposition

$$v = \sum_{\alpha=1}^{r} a_{\alpha_1,\alpha_2} \otimes a_{\alpha_2,\alpha_3} \otimes \ldots \otimes a_{\alpha_n,\alpha_1} \tag{1.6}$$

is C_n invariant, as any cyclic permutation leaves it unchanged. More generally, in a G invariant decomposition, the term $v_\beta^{[i]}$ at site i must equal the term $v_{g^{-1}\beta}^{[gi]}$ at site gi, where the summation index has changed from β to $g^{-1}\beta$, and the part index has also changed from i to gi. If all elements in the decomposition satisfy this property, we say that the parts are G invariant. This gives rise to an (Ω, G)-decomposition (whose formal definition can be found in page 10 of P1). The bottom line is that, in an (Ω, G)-decomposition, the elementary tensors satisfy some symmetry conditions.

Now, by construction, if v admits an (Ω, G)-decomposition, then v is G invariant. Q2 (and the main question considered in P1) is the reverse one:

If v is G-invariant, does it admit an (Ω, G) decomposition?

The answer is: Yes, if G acts freely on Ω (Theorem 13 of P1). Moreover, the multiplicity of the facets of Ω can always be increased to satisfy this condition (Proposition 7 of P1). In other words: we give a sufficient condition for the existence of an (Ω, G)-decomposition, namely the freeness of the action of G on Ω, and show that this condition can always be satisfied by increasing the multiplicity of the facets. (Freeness is given in Definition 4 of P1).

Our result generalises the symmetric tensor rank decomposition (considered in mathematics) or the translationally invariant operator Schmidt decompositions (used to describe chains of quantum spins). It also allows to compare the (Ω, G) ranks if the group changes (Section 4.2 of P1) or if the simplicial complex changes (Section 4.3 of P1). For example, we show that the symmetric tensor rank is the largest of the (Ω, G) ranks. Additionally, it lets us address Q3 by combining invariance and positivity, as we will now see.

1.3 If the whole is positive, are the parts positive too?

Analogously to Q2, let us start by asking: What does it mean that the whole is positive? It means that v is in a certain cone—typically, the cone of positive semidefinite matrices, or the cone of nonnegative tensors. And what does it mean that the parts are positive? It can mean three things:

(i) In the *separable* case, it means that the parts are in the same cone as v.

(ii) In the *purification* case, it consists of a decomposition of a certificate that v is in that cone. For example, if v is in the cone of the positive semidefinite matrices, the purification is a decomposition of A, where $v = A^{\dagger}A$.[3]

(iii) In the *quantum square root* case, it consists of a decomposition of a *Hermitian* certificate that v is in that cone. For example, if v is in the cone of the positive semidefinite matrices, the quantum square root is a decomposition of A, where $v = A^2$.

More formally:

(i) A *separable* (Ω, G) *decomposition* is an (Ω, G) decomposition where every $v_{\beta}^{[i]}$ is positive semidefinite. The minimal number of terms is called the (Ω, G) separable rank, denoted $\text{sep-rank}_{(\Omega,G)}$.

(ii) An (Ω, G) *purification* of v is an (Ω, G) decomposition of L where $v = L^{\dagger}L$. The minimal number of terms is called the (Ω, G) purification rank, denoted $\text{puri-rank}_{(\Omega,G)}$.

(iii) An (Ω, G) *quantum square root* of v is an (Ω, G) decomposition of L where $v = L^2$. The minimal number of terms is called the (Ω, G) quantum square root rank, denoted $\text{q-sqrt-rank}_{(\Omega,G)}$.

[3]For the mathematician: † denotes complex conjugate transpose, in mathematics usually denoted *.

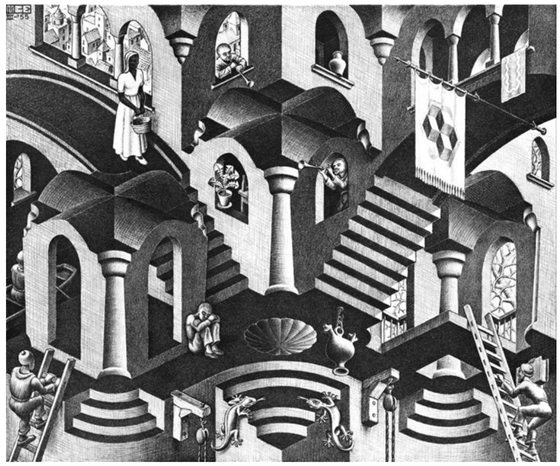

Fig. 1.2: Positivity gives rise to the notion of convexity, which interacts with the multiplicity of systems in a surprising way. (Lithograph by M. C. Escher).

Now we can address Q3, together with invariance:

If v is G-invariant and positive, when does it admit each of the three positive (Ω, G) decompositions?

The short answer is that v will generally admit each of the three positive invariant decompositions, *but generally for a very large prize*. That is, there are *separations* among each of these ranks. We know this because in the very special case in which v just has two parts, the above decompositions correspond to the nonnegative factorisation, the psd factorisation and the square root factorisation, respectively, and there are separations among these already. In plain words, what happens is that cones interact in a very rich way with the multiplicity of systems (Fig. 1.2), and this is already the case for the classical (i.e. commutative) and bipartite case. Let us explain this in more detail for the bipartite case.

Correspondence with matrix decompositions

Our framework, and in particular our three definitions of positivity of page 8, generalise other well-known decompositions. Specifically, in the simplest non-trivial case, when the simplicial complex Ω consists of two vertices sharing an edge, our decompositions specialise to the well-studied non-negative [CR93], positive semidefinite [FGP+15], completely positive [BSM03] and completely positive semidefinite transposed decompositions P2 of matrices. Let us explain this.

For matrices, there are two main notions of positivity. The first is that M be *nonnegative*, meaning entrywise nonnegative. This notion of positivity (or nonnegativity) is in essence the same as that of a nonnegative vector. The second is that M be *positive semidefinite*, namely diagonalisable and with nonnegative eigenvalues. M could thus be real and symmetric, or Hermitian—in either case with nonnegative eigenvalues. For the quantum case the latter is the important one. This notion of positivity is inherent to a matrix—the matrix itself can have complex entries, but its eigenvalues must be nonnegative.

For a matrix M, Q3 becomes:

If M has some notion of positivity, can M be decomposed so that it preserves this notion of positivity?

If M is nonnegative, the *nonnegative factorisation* is defined as

$$M = AB \qquad \text{where } A \text{ and } B \text{ are nonnegative} \tag{1.7}$$

and the *nonnegative rank* is the minimal number of columns of A, denoted nn-rank (or rank_+). A noncommutative version of the nonnegative factorisation is the *positive semidefinite (psd) factorisation*, which is defined as

$$M_{i,j} = \text{tr}(A_i B_j) \quad \text{where all } A_i \text{ and } B_j \text{ are positive semidefinite.} \tag{1.8}$$

The *psd rank* is defined as the minimal size of all A_is and B_js that satisfy (1.8). (Note that there need to be as many A_is (B_js) as the number of rows (columns) of M, so this cannot define a rank.) Usually these psd matrices are defined in the reals [FMP+12, FGP+15] (i.e. they are real symmetric matrices with nonnegative eigenvalues), although for our connection the complex case is the relevant one. Since in the psd factorisation we could choose all A_i and B_j to be diagonal, and we would recover a nonnegative factorisation, it follows that it is harder to decompose with nonnegative numbers than with real numbers (rank ≤ nn-rank), that noncommutativity helps (psd-rank ≤ nn-rank), and that it does not get smaller than the rank.[4]

Why is this interesting? Because *the nonnegative rank and the psd rank are much more expensive than the rank*. In plain words, negative numbers allow for massive shortcuts in a finite set of sums (even if the result of these sums is positive). Formally, there is a *separation* between each of these ranks: there is a sequence of matrices M_n (whose size increases with n) such that $\text{rank}(M_n)$ is bounded, but

[4]The precise inequality is $\frac{1}{2}\sqrt{1+8\text{rank}(M)} - \frac{1}{2} \leq \text{psd-rank}_{\mathbb{R}}(M) \leq \text{nn-rank}(M)$.

$$\begin{tikzcd}
\text{minimal factorisation} \arrow[r, "\text{symmetric}"] \arrow[d, "\text{nonnegative}"'] & \text{symmetric factorisation} \arrow[d, "\text{nonnegative}"] \\
\text{nonnegative factorisation} \arrow[r, "\text{symmetric}"] \arrow[d, "\text{non-commutative}"'] & \text{cp factorisation} \arrow[d, "\text{non-commutative}"] \\
\text{psd factorisation} \arrow[r, "\text{symmetric}"] & \text{cpsd factorisation}
\end{tikzcd}$$

Fig. 1.3: Existing factorisations of matrices, including symmetry (on the right column). In the top row, the factorisation does not impose any positivity condition on the local terms. In the middle row, the factorisation demands that the local matrices be nonnegative, whereas in the bottom row, it asks for a non-commutative version of nonnegativity, namely that of Eq. (1.8). Each factorisation has an associated rank; that of the minimal factorisation is the usual rank.

nn-rank(M_n) diverges. This means that rank cannot be upper bounded by a function of nn-rank exclusively, and we write rank $\ll$ nn-rank. The same is true for the rank and psd rank, rank $\ll$ psd-rank, and for the psd rank and nonnegative rank, psd-rank $\ll$ nn-rank [GPT13]. So there are separations everywhere.

Let us now add symmetry of M as a further ingredient, which for M a matrix just means that M be symmetric (i.e. $M = M^t$ if real, and $M = M^\dagger$ if complex). The three decompositions above in the symmetric case are, first, the *symmetric factorisation*, defined as

$$M = AA^t \qquad \text{where } A \text{ is complex} \tag{1.9}$$

where the minimal number of columns of A is the symmetric rank. Second, the *cp factorisation* (standing for completely positive), defined as

$$M = AA^t \qquad \text{where } A \text{ is nonnegative} \tag{1.10}$$

and the minimal number of columns of A is the cp rank. And third, the *cpsd factorisation* (standing for completely positive semidefinite), defined as

$$M_{i,j} = \mathrm{tr}(A_i A_j) \qquad \text{where } A_i \text{ is positive semidefinite} \tag{1.11}$$

and the minimal size of all A_i's is the cpsd rank (see Fig. 1.3). So 'completely' here means 'symmetric'.

Now, our framework of (Ω, G) decompositions, with the notions of positivity of page 8, provides a non-commutative generalisation of the six decompositions of Fig. 1.3. To see this, we now consider M to be a bipartite operator (instead of a bipartite matrix), i.e. $M \in \mathcal{M}_{d_1} \otimes \mathcal{M}_{d_2}$, where $\mathcal{M}_d$ is the space of $d \times d$ complex matrices. The *operator Schmidt decomposition* of M is defined as

$$M = \sum_{\alpha=1}^{r} A_\alpha \otimes B_\alpha \tag{1.12}$$

and the minimal such r is the operator Schmidt rank. The *separable decomposition* of M is defined as

$$M = \sum_{\alpha=1}^{r} \sigma_\alpha \otimes \tau_\alpha \quad \text{where } \sigma_\alpha, \tau_\alpha \text{ are positive semidefinite} \tag{1.13}$$

where the minimal such r is the *separable rank*. (This only exists if M is in the convex cone of the Cartesian product of the cones of positive semidefinite matrices, i.e. if M is separable). The *local purification* of M is defined as

$$M = LL^\dagger, \qquad \text{where } L = \sum_{\alpha=1}^{r} A_\alpha \otimes B_\alpha \tag{1.14}$$

where L need not be a square matrix,[5] and the minimal such r is the *purification rank*. The *quantum square root* of M is

$$M = L^2 \qquad \text{where } L = \sum_{\alpha=1}^{r} A_\alpha \otimes B_\alpha \tag{1.15}$$

and the *q-sqrt-rank* is the operator Schmidt rank of L.

If M is additionally symmetric, i.e. invariant under the permutation of parts 1 and 2, we have the *t.i. operator Schmidt decomposition* (where t.i. stands for translationally invariant), which is given by

$$M = \sum_{\alpha=1}^{r} A_\alpha \otimes A_\alpha \tag{1.16}$$

and the minimal such r is the *t.i. operator Schmidt rank*. The *t.i. separable decomposition* is given by

$$M = \sum_{\alpha=1}^{r} \sigma_\alpha \otimes \sigma_\alpha \qquad \text{where } \sigma_\alpha \text{ is positive semidefinite} \tag{1.17}$$

[5]For the physicists: this the same as a purification $M = \mathrm{tr}_{\mathrm{aux}}|\psi\rangle\langle\psi|$. The sum over auxiliary indices aux becomes the internal sum in the matrix multiplication of $LL^\dagger$ in (1.14).

$$\begin{tikzcd}
\text{operator Schmidt decomposition} \arrow[r, "\text{symmetric}"] \arrow[d, "\text{nonnegative}"'] & \text{t.i. operator Schmidt decomposition} \arrow[d, "\text{nonnegative}"] \\
\text{separable decomposition} \arrow[r, "\text{symmetric}"] \arrow[d, "\text{non-commutative}"'] & \text{t.i. separable decomposition} \arrow[d, "\text{non-commutative}"] \\
\text{local purification} \arrow[r, "\text{symmetric}"] & \text{t.i. local purification}
\end{tikzcd}$$

Fig. 1.4: Our non-commutative generalisation of Fig. 1.3, concerning factorisations of bipartite operators (instead of matrices), without and with symmetry (left versus right hand side). The top row contains the factorisation without any certificate of positivity in the parts. The middle row contains the separable decomposition, which essentially imposes nonnegativity. The bottom row contains a non-commutative version of positivity, namely a purification, where the parts contain a certificate of positivity in a non-trivial way. Just one detail is not matching: the t.i. local purification generalises the cpsd factorisation with an additional transpose, which we term the cpsd*t* decomposition (see P5 and P1 for details).

and the minimal such r is the *ti-separable rank*. And the *t.i. local purification* is

$$M = AA^\dagger \qquad \text{where } A = \sum_{\alpha=1}^{r} B_\alpha \otimes B_\alpha \tag{1.18}$$

and the minimal such r is the *t.i. purification rank*. These six decompositions for a bipartite operator M are summarised in Fig. 1.4.

What we show in P1 (and partially in P5) is that, if M is a bipartite operator *which is diagonal in the computational basis*, then Fig. 1.4 becomes Fig. 1.3. In words, this says that our generalisation is sensible. More formally, if $M = \text{diag}(N)$ where N is a nonnegative matrix and $\text{diag}(N)$ rearranges the entries of N into a diagonal, the factorisations of the psd matrix M coincide with the factorisations of the nonnegative matrix N, up to an extra transpose in the cpsd factorisation (giving rise to the cpsdt factorisation; see Table 1.1).

Since there is a separation between rank, nonnegative rank and psd rank (i.e. the left column of Fig. 1.3), from the correspondence of Table 1.1 it follows that there is a separation between the operator Schmidt rank, the purification rank and the separable rank (i.e. the left column of Fig. 1.4). In symbols, osr $\ll$ puri-rank $\ll$ sep-rank. (The former separation was a central result of [DSPGC13]).

Finally, the correspondence of Table 1.1 also inspired us to generalise the results for matrices rank two (explained, e.g., in [FGP+15]) to operators with operator Schmidt rank two in P4, summarised in Table 1.2. This hopefully illustrates the usefulness of our generalisation.

Decomposition of $M = \text{diag}(N)$	Decomposition of N
operator Schmidt decomposition	minimal factorisation
separable decomposition	nonnegative factorisation
local purification	complex psd factorisation
t.i. operator Schmidt decomposition	symmetric factorisation
t.i. separable decomposition	cp factorisation
t.i. local purification	complex cpsdt factorisation

Tab. 1.1: If a psd matrix M is diagonal in the computational basis, $M = \text{diag}(N)$ where N is the nonnegative matrix containing the diagonal of M, then the decompositions of M on the left hand side are the same as the decompositions of N on the right hand side P5, except for an additional transpose in the cpsd factorisation.

Nonnegative matrix M	rank(M)	
	1	Trivial (all ranks the same)
	2	nn-rank = psd-rank = 2
	3	nn-rank and psd-rank can be unbounded

Positive semidefinite matrix ρ	osr(ρ)	
	1	Product state (all ranks the same)
	2	Separable, and sep-rank = puri-rank = 2
	3	puri-rank and sep-rank can be unbounded

Tab. 1.2: The case of rank 1 is trivial, of rank 2 is easy and fully characterised, and of rank 3 is as hard as the general case, i.e. it already shows separations. This is true both for a nonnegative matrix M [FGP$^+$15] and for a bipartite positive semidefinite matrix ρ, as we showed in P4.

1.4 And in the approximate case?

Does the answer to Q3 change if the parts only need to reproduce the whole up to some ε? P2 studies (Ω, G) decompositions with positivity in the *approximate* case (Fig. 1.5), where the notion of approximation is given by ε-balls around the elements, in some norm (Schatten p-norm or ℓ_p norm). For example, we define

$$\text{rank}^{\varepsilon}_{(\Omega,G)}(M) = \min_{N \in B_\varepsilon(M)} \text{rank}_{(\Omega,G)}(N), \tag{1.19}$$

where $B_\varepsilon(M)$ is the ball around M measured in the given norm, and similarly for the other cases, giving rise to sep-rank$^{\varepsilon}_{(\Omega,G)}$, puri-rank$^{\varepsilon}_{(\Omega,G)}$ and q-sqrt-rank$^{\varepsilon}_{(\Omega,G)}$. In fact every such approximate rank has an additional superindex p indicating the Schatten p-norm in which the distance in the ball B_ε is measured, or ℓ_p if the distance is measured in ℓ_p norm, as the upper bounds on these approximate ranks will depend on the norm.

The central result of P2 is that essentially all separations among ranks disappear

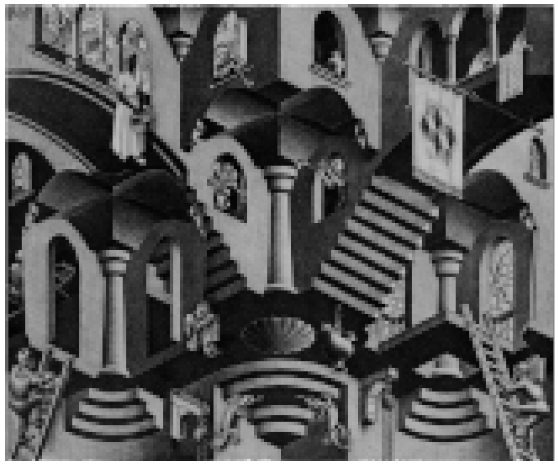

Fig. 1.5: The interaction of positivity (convexity) and the multiplicity of systems with a finite resolution (i.e. with ε balls) is very different than in the exact case, as many separations between ranks disappear, as we showed in P2. So the answer to Q4 is very different to that of Q3. (Compare with Fig. 1.2).

in the approximate case (Corollary 26 of P2). In other words, if the parts need only represent the whole approximately, the cost can be much lower. To prove the result, we leverage the recent *approximate Carathéodory Theorem* [Iva19], which shows that every element in a convex hull can be approximately represented by using a number of extremal points which is independent of the dimension of the ambient vector space, and only depends on the error. This leads to our central technical result (Theorem 22 of P2).

It is worth mentioning two caveats to the general results of P2. The first is that our results only hold for certain ranges of norms, namely $1 < p \leq 4/3$, $p = 2$ or $4 \leq p < \infty$. We do not know if they can be extended to the missing ranges of p. In particular, at the moment they do not hold for $p = 1$, i.e. the trace norm, which is the relevant one for quantum mechanics. The second caveat is that the upper bounds on the ε ranks often involve a "gauge function" (called $\mu_{\sqrt{\rho}}$ in P2), which is needed in order to consider a convex set, instead of just a convex cone, and which is difficult to compute. Moreover, our results do not exploit the tensor structure of v (they only use that v be in a cone, and then the general tensor rank decomposition), so there is room for improvement.

1.5 What about other worlds with the same parts–whole relation?

While the framework of (Ω, G)-decompositions is inspired by tensors—in particular, by the description of quantum many-body systems—, it applies to any tensor product structure. In P3, we apply it to *real multivariate polynomials*. These are objects in the tensor product space of polynomials in each of their variables,

$$\mathscr{P} := \mathbb{R}[\mathbf{x}^{[1]}, \mathbf{x}^{[2]}, \ldots, \mathbf{x}^{[n]}] \cong \mathbb{R}[\mathbf{x}^{[1]}] \otimes \mathbb{R}[\mathbf{x}^{[2]}] \otimes \cdots \otimes \mathbb{R}[\mathbf{x}^{[n]}],$$

where $\otimes$ denotes the algebraic tensor product and $\mathbf{x}^{[i]}$ a collection of variables $x_1^{[i]}, \ldots x_{m_i}^{[i]}$. We translate Q1, Q2 and Q3 to, respectively:

(1) Every polynomial $p \in \mathscr{P}$ can be expressed as a finite sum of "elementary constituents"

$$p^{[1]}(\mathbf{x}^{[1]}) \cdot p^{[2]}(\mathbf{x}^{[2]}) \cdots p^{[n]}(\mathbf{x}^{[n]}),$$

but how are the summation indices arranged?

(2) If p is symmetric under the exchange of, say, systems $[i]$ and $[j]$, when can this symmetry be reflected in the decomposition, and at what price?

(3) If p is positive (for some notion of positivity), when can this positivity be reflected in the decomposition, and at what price?

Our framework solves these three questions in the following way—in particular applied to polynomials:

(1) The summation structure is described by a weighted simplicial complex Ω, so that every system i is associated to a vertex of Ω, and every summation index to a facet of Ω. For example, the indices can be arranged in a circle,

$$p = \sum_{\alpha_1, \ldots, \alpha_n = 1}^{r} p^{[1]}_{\alpha_1, \alpha_2}(\mathbf{x}^{[1]}) \cdot p^{[2]}_{\alpha_2, \alpha_3}(\mathbf{x}^{[2]}) \cdots p^{[n]}_{\alpha_n, \alpha_1}(\mathbf{x}^{[n]}). \tag{1.20}$$

or with a single index,

$$p = \sum_{\alpha=1}^{r} p^{[1]}_{\alpha}(\mathbf{x}^{[1]}) \cdot p^{[2]}_{\alpha}(\mathbf{x}^{[2]}) \cdots p^{[n]}_{\alpha}(\mathbf{x}^{[n]}). \tag{1.21}$$

(2) By definition, an (Ω, G)-decomposition of a polynomial contains a certificate of invariance under the group G. We characterise which G-invariant polynomials admit an (Ω, G)-decomposition.

Our framework models symmetries as follows: we have a group G acting on the set $\{1,\dots,n\}$, and the induced action on the polynomial space $\mathscr{P}$ is obtained by permuting system $[i]$ to $[gi]$,

$$g : \mathbf{x}^{[i]} \mapsto g\mathbf{x}^{[i]} := \mathbf{x}^{[gi]}.$$

A polynomial is G-invariant if it is invariant with respect to all such permutations $g \in G$, and we want to make this invariance explicit in the decomposition of p. For example, the following decomposition

$$p = \sum_{\alpha_1,\dots,\alpha_n=1}^{r} p_{\alpha_1,\alpha_2}(\mathbf{x}^{[1]}) \cdot p_{\alpha_2,\alpha_3}(\mathbf{x}^{[2]}) \cdots p_{\alpha_n,\alpha_1}(\mathbf{x}^{[n]})$$

makes explicit that p is invariant under the cyclic group, $\mathbf{x}^{[i]} \mapsto \mathbf{x}^{[i+1]}$, and

$$p = \sum_{\alpha=1}^{r} p_\alpha(\mathbf{x}^{[1]}) \cdot p_\alpha(\mathbf{x}^{[2]}) \cdots p_\alpha(\mathbf{x}^{[n]})$$

makes explicit that p is invariant under the full symmetry group. In our framework, the former corresponds to the circle with the cyclic group, and the minimal number of terms is the *t.i. operator Schmidt rank*, and the latter to the *symmetric tensor decomposition*, and the minimal number of terms is the *symmetric tensor rank* [CGLM08, Shi18] (cf. (2)).

(3) By definition, a separable or sum-of-squares (sos) (Ω, G)-decomposition contains a certificate of invariance and of membership in the separable or sos cone, respectively. We characterise which separable or sos polynomials admit such decompositions.

Our contributions to the description of multivariate polynomials can be summarised as follows:

- We define an (Ω, G) decomposition of a polynomial, as well as separable- and sum of squares (Ω, G) decomposition, provide existence theorems of invariant decompositions (Theorem 32 of P3), inequalities between the ranks (Proposition 39) and separations between them (Corollary 45 of P3).
- We show that the rank separations disappear in the approximate case (Theorem 49 of P3).
- We show that a polynomial problem is undecidable by linking it to a positivity problem of Matrix Product Operators [DCC+16] (Theorem 51 of P3).

It is worth mentioning that we are currently studying approximate tensor decompositions from the perspective of *border ranks* [DKN22], where tensor decompositions are again considered within an approximation error $\varepsilon > 0$, which is however not fixed but tends to 0. Naively, one would expect that arbitrarily close approximations become (sooner or later) equal to the exact description. While this is true for many matrix factorizations (including the standard matrix rank or positive factorizations [FGP$^+$15]), it is false for tensors [LQY11, CVZ19], leading to the so-called border rank phenomenon. We are currently generalizing several border rank results to Ω factorisations, also including invariance and positivity.

1.6 Given some shadows of the whole, is the whole positive?

In many cases, the whole is unassailable, because it lives in an exponentially big space. In this context, it is natural to assume that we only have access to some shadows thereof, or 'traces' of its existence, which in our case means a few moments. The question is then:

Given a few moments of a Hermitian matrix, what can we assert about its positivity?

More specifically, can we say at least and at most how far it is from the cone of positive semidefinite matrices?

This question is natural if the whole has a tensor structure with few matrices in each part, that is, if some of its Ω ranks is small, for in this case the moments are easily computed whereas a diagonalisation of M is out of reach.

This is the question we embark to investigate in P6. This work sheds some light on the interaction between the cones, i.e. the positivity structure, and the multiplicity of systems—the 'universes' of the Prologue.

Let us start by motivating the assumption that we only be able to see some 'shadows' of the whole. The problem starts with the simple but far-reaching observation that the dimension of $V = V_1 \otimes \ldots \otimes V_n$ be the product of the dimensions of each V_i, so it grows exponentially with n. This is dramatic for the description of quantum many-body systems, and is the origin of the program of tensor networks [CPGSV21, Oru19], whose goal is to develop scalable descriptions of these systems. In the tensor network paradigm, it is natural to use a few matrices for each local Hilbert space V_i. For example, the state of the system in one spatial dimension is described by

$$M = \sum_{j_1,\ldots,j_n=1}^{r} A^{[1]}_{j_1} \otimes A^{[2]}_{j_1,j_2} \otimes \cdots \otimes A^{[n]}_{j_{n-1}}$$

where each local matrix is Hermitian (this is the decomposition of (1.2)), but not necessarily positive semidefinite. Yet, M must be positive semidefinite to describe a quantum state. While there is a way of imposing positivity in the local matrices, resulting in the local purification form (page 8), this generally comes at the price of a very large increase in the number of matrices, i.e. there is a separation among the two decompositions [DSPGC13], as we explained in Section 1.3. In the case of arbitrarily many parts (i.e. arbitrarily large n), the problem becomes undecidable [DCC+16], and its bounded version NP-hard [KGE14].

In the context of free spectrahedra, it is also natural to assume that we have access to a few moments of a large Hermitian matrix (cf. page iv and P6). That is, we assume that M can be expressed as

$$M = \sum_{j=1}^{r} A_j^{[1]} \otimes \cdots \otimes A_j^{[n]}$$

where each $A_j^{[i]}$ is Hermitian and of reasonably small dimension. (Note that this is the tensor rank decomposition of (1.1).) Note also that the relation between positivity and the multiplicity of systems is very simple if $r = 1$: M is positive semidefinite if and only if each local matrix $A^{[i]}$ is either positive semidefinite or negative semidefinite, and the number of negative semidefinite matrices is even. But for $r > 1$ it is unclear whether any simple criterion exists at all, and the separations explained in Section 1.3 suggest that this is not the case—instead, we should imagine a rich landscape between the cones and the universes, as that of Fig. 1.2.

In summary, in several contexts the whole is a Hermitian operator M on a Hilbert space whose dimension is so large that it is impossible to write down all matrix entries in an orthonormal basis. Moreover, it is natural to assume that we only have access to a few moments of M. How does one determine whether such M is positive semidefinite? In P6 we approach this problem by deriving asymptotically optimal bounds to the distance to the positive semidefinite cone in Schatten p-norm for all integer $p \in [1, \infty)$, assuming that we know the moments $\mathrm{tr}(M^k)$ up to a certain order $k = 1, \dots, m$.

More specifically, we address the following questions:

(i) Given the first m normalized moments

$$\mathrm{tr}(M^k) = \frac{1}{s}\mathrm{tr}(M)$$

where s is the size of the matrix, for $k = 1, \dots, m$, of a Hermitian operator M with $||M||_\infty \leq 1$, can one show that M is not positive semidefinite?

(ii) Given these moments and a $p \in [1, \infty)$, can one optimally bound the distance of M to the positive semidefinite cone from above and below in Schatten p-norm?

Since both the moments and the positive semidefiniteness of a Hermitian operator M are characterized by the distribution of eigenvalues, we are secretly concerned with a version of the truncated Hausdorff moment problem, as explained in P6.

The main idea to address these questions is, first, to note that the distance to the cone of positive semidefinite matrices is given by the norm of the negative part of M, and, second, to approximate this norm with the few moments. Namely, a Hermitian M can be expressed as

$$M = M_+ - M_-$$

where $M_- = f_1(M)$, where $f_1(x)$ is the absolute value of the negative part of x, and 0 elsewhere, and M_+ is the difference between M and M_-. The distance from M to the cone of positive semidefinite matrices (in Schatten p-norm) is, by definition,

$$d_p(M) = \inf_{N \geqslant 0} ||M - N||_p$$

where $N \geqslant 0$ denotes that N be positive semidefinite. This distance is given precisely by the norm of the negative part of M (Lemma 1 of P6), namely

$$d_p(M) = ||M_-||_p.$$

So, ideally, we would like to compute $d_p(M)$, but we only have access to a few moments of M, so we instead estimate $d_p(M)$ as accurately as possible. Namely, we consider the best upper and lower bounds to $d_p(M)$ that can be obtained from these moments, and we provide three methods to compute these bounds and relaxations thereof:

The sos polynomial method computes the upper and lower bounds to $d_p(M)$ by solving a semidefinite program.

The Handelman method is a linear program relaxation, which uses another ansatz for nonnegative polynomials on $[0, 1]$ to compute the upper and lower bounds to $d_p(M)$.

The Chebyshev method is a relaxation not involving any optimization. The polynomials are the Chebyshev polynomials (whose degree is at most m) that best approximate the negative part function.

In P6 we investigate the analytical and numerical performance of these three methods and present a number of example computations.

1.7 Which wholes can be divided into arbitrarily many parts?

This question means: Which states have a continuum limit? In P7, P8 and P9 we address this question for Matrix Product States [PGVWC07, Vid03, FNW92] in the following way.

We consider Matrix Product States (MPS), more particularly for families of translationally invariant MPS. These correspond to families of tensor decompositions where, in the language of the previous questions, the simplicial complex is the circle graph and the symmetry group the cyclic group C_n. Namely, a family $\mathcal{V}(A)$ is defined as

$$\mathcal{V}(A) = \{|V_N(A)\rangle\}_{N\in\mathbb{N}}$$

where

$$|V_N(A)\rangle = \sum_{i_1,\ldots,i_N=1}^{d} \mathrm{tr}(A^{i_1}\cdots A^{i_N})|i_1,\ldots,i_N\rangle$$

where $A = \{A^i_{\alpha,\beta}\}$ with $i = 1,\ldots,d$ and $\alpha,\beta = 1,\ldots D$ is a tensor with complex entries. The goal is to characterise which families $\mathcal{V}(A)$ have a continuum limit in terms of the mathematical properties of the tensor A. Moreover, if $\mathcal{V}(A)$ has a continuum limit, we also want to specify it. Note that A provides a 0-dimensional characterisation of the family, as it does not scale with the system size—we can think of A as a *grammar* of $\mathcal{V}(A)$. (A grammar in the sense of formal languages, see e.g. [SDD20] or [Koz97]).

In the pursuit of this goal, a crucial tool is the canonical form of MPS and its associated fundamental theorem [CPGSV17], which relates different MPS representations of a state. In words, this theorem says that, if A and B are in canonical form, $\mathcal{V}(A) = \mathcal{V}(B)$ if and only if A and B are related by a basis change common to all physical indices. The crucial part of the theorem is the "only if", for it allows to go from a global condition ($\mathcal{V}(A) = \mathcal{V}(B)$) to a local condition (how A and B are related). This theorem fully characterises the freedom of the map $A \to \mathcal{V}(A)$, and gives us a firm handle on the characterisation of the family $\mathcal{V}(A)$ in terms of A. For this reason, the canonical form and the fundamental theorem underpin many of the analytical results derived through MPS, such as the celebrated classification of symmetry-protected phases in one dimension [SPGC11] (see also [CGW11]).

Yet, the canonical form is only defined for MPS without non-trivial periods. Specifically, it excludes translationally invariant MPS which are superpositions of states with non-trivial periodicity $m > 1$, such as the paradigmatic antiferromagnetic state

$$|0,1,0,1,\ldots\rangle + |1,0,1,0,\ldots\rangle$$

with period $m = 2$. Note that, after blocking m spins we obtain a translationally invariant MPS in a trivial way, at which point the canonical form and the fundamental theorem can be applied. However, the local entanglement structure relating to the non-trivial periodicity is lost in this procedure, and the physical properties of the system can change radically (e.g., the antiferromagnet becomes a ferromagnet when blocking 2 sites). More importantly for our purposes, the investigation of which states can be fine-grained—the stepping stone on which the continuum limit will be defined—does not allow for blocking of the sites.

In P7 we introduce a new standard form for MPS, the irreducible form, which is defined for arbitrary MPS, including periodic states. We show that any tensor can be transformed into a tensor in irreducible form describing the same MPS. We then prove a fundamental theorem for MPS in irreducible form. In words, it says that if two tensors in irreducible form give rise to the same MPS, they must be related by a similarity transform, together with a matrix of phases. Slightly more formally, given any A and B in irreducible form, $\mathcal{V}(A) = \mathcal{V}(B)$ if and only if there is a unitary Z and an invertible matrix Y such that $ZA^i = YB^iY^{-1}$ for all i, with $[Z, A^i] = 0$ and $\mathcal{V}(A) = \mathcal{V}(ZA)$.

As an application of this result, we provide an equivalence between the refinement properties of a state and the divisibility properties of its transfer matrix. This hinges on the very fruitful connection between MPS and quantum channels, i.e. trace-preserving completely positive maps. This stems from the observation that, given a tensor A, its transfer matrix

$$E_A = \sum_{i=1}^{d} A^i \otimes \bar{A}^i \tag{1.22}$$

(where $\bar{\ }$ denotes complex conjugate) is a matrix representation of a completely positive map, whose Kraus operators are precisely A^i. This map can be made trace-preserving too. The theory of quantum channels (notably exposed in the unpublished notes [Wol11]) can be used to gain much insight into properties of A, as we in fact do in P7.

In particular, given a tensor B, we say that $\mathcal{V}(B)$ can be p-refined if there exists another tensor A and an isometry $W : \mathbb{C}^d \to (\mathbb{C}^d)^{\otimes p}$ such that

$$|V_{pN}(A)\rangle = W^{\otimes N}|V_N(B)\rangle \quad \forall N.$$

On the other hand, a trace-preserving completely positive map E is called p-divisible if there is another trace-preserving completely map T such that $E = T^p$, where the latter denotes the p-fold application of the map (both definitions are provided in P8). We show in P7 that if B is in irreducible form then these two concepts are the same, i.e. $\mathcal{V}(B)$ can be p-refined if and only if its transfer matrix E_B is p-divisible (Theorem 17 of P7).

In P8 we determine which translationally invariant matrix product states have a continuum limit, that is, which can be considered as discretized versions of states defined in the continuum. Specifically, we say that $\mathcal{V}(B)$ has a continuum limit if there is a $p > 1$ such that the procedure of p-refining ℓ times followed by the blocking of $n_\ell \in \mathbb{N}$ of the resulting spins converges in ℓ, as long as $(n_\ell/p^\ell)_\ell \to 0$ as $\ell \to \infty$. Here $(n_\ell/p^\ell)_\ell$ denotes the infinite sequence whose elements are n_ℓ/p^ℓ with $\ell \in \mathbb{N}$. The main result of P8 (Theorem 1) is that, given $\mathcal{V}(B)$ with B in irreducible form, the following statements are equivalent:

1. $\mathcal{V}(B)$ has a continuum limit.

2. E_B is infinitely divisible.

3. There is a quantum channel P and a Liouvillian of Lindblad form L such that $E_B = Pe^L$, $P^2 = P$, and $PLP = PL$.

Our contribution is the equivalence of 1. and 2., as the equivalence of 2. and 3. was proven by Denisov [Den89] and Kholevo (usually spelled Holevo) [Kho87]. If the projector is trivial, $P = I$, the corresponding transfer matrix e^L coincides with that of a translationally invariant continuous MPS [VC10, HCOV13], but if it is not, $P \neq I$, the continuum limit cannot be represented by a continuous MPS. We will address this question in P9.

We also say that $\mathcal{V}(A)$ has a coarse continuum limit if there is a $\mathcal{V}(B)$ and an $n \in \mathbb{N}$ such that $\mathcal{V}(A)$ is the n-refinement of $\mathcal{V}(B)$, and $\mathcal{V}(B)$ has a continuum limit. It follows from the main theorem that $\mathcal{V}(A)$ has a coarse continuum limit if and only if there exists an $n \in \mathbb{N}$ such that E_A^n is infinitely divisible. The antiferromagnet, the one-dimensional cluster state [RB01], and the AKLT state [AKLT87] are examples of states with a coarse continuum limit but no continuum limit.

One surprising outcome of the characterisation of P8 is precisely the fact that continuous MPS (cMPS) do not capture all continuum limits of MPS. The missing element is the projector P in the transfer matrix. In P9 we provide a generalized ansatz of cMPS that is capable of expressing the continuum limit of any MPS. It consists of a sum of cMPS with different boundary conditions, each attached to an ancilla state. This ansatz can be interpreted as the concatenation of a state which is at the closure of the set of cMPS together with a standard cMPS. The first can be seen as a cMPS in the thermodynamic limit, or with matrices of unbounded norm. We provide several examples and discuss the result.

1.8 Which composition rules preserve the order of the parts?

If the parts are in a cone, they build a partial order (usually denoted with the symbol $\leqslant$), which is ultimately inherited from the order of real numbers. Yet, in composing the parts, this order is generally not preserved. So we consider the question: Which composition rules preserve the order relation of the parts? This question is relevant for recent approaches to (quantum) natural language processing, which represent the meaning of words by positive semidefinite matrices, and hyponymy by the $\leqslant$ relation, as we will now see.

Positive semidefinite matrices have a partial order defined on them, sometimes called the Löwner order, so that

$$A \geqslant B \quad \text{if and only if} \quad A - B \geqslant 0$$

where the latter symbol means positive semidefinite (and we slightly abuse notation by using the same symbol for $A \geqslant B$). We also write $C \leqslant 0$ if and only if $-C \geqslant 0$. In some approaches to compositional distributional semantics, the meaning of a word is represented by a positive semidefinite matrix [BCLM18], and the Löwner order is used to represent the hyponymy relation. The hyponomy relation is the relation of entailment—for example, a cat is a mammal, so cat is a hyponym of mammal (and mammal is a hypernym of cat). Similarly, for verbs, climb entails move, so climb is a hyponym of move. This relation is represented with the Löwner order, so that e.g.

$$[\![\text{cat}]\!] \leqslant [\![\text{mammal}]\!]$$

where $[\![\text{cat}]\!]$ denotes the positive semidefinite matrix representing the meaning of cat. Similarly, $[\![\text{climb}]\!] \leqslant [\![\text{move}]\!]$.

In order to obtain the meaning of phrases and sentences, the meaning of its words is combined, for example, with composition rules like Fuzz and Phaser [CM20] (a.k.a. KMult and BMult [Lew19]; see [CSC10] for a background to this approach, and [PKCS15] for the extension to positive semidefinite matrices). Yet, these rules do not preserve hyponymy, that is, from

$$[\![\text{cat}]\!] \leqslant [\![\text{mammal}]\!] \quad \text{and} \quad [\![\text{climb}]\!] \leqslant [\![\text{move}]\!]$$

it does not follow that

$$[\![\text{cats climb}]\!] \leqslant [\![\text{mammals move}]\!].$$

In P10, we introduce a generic way of composing the positive semidefinite matrices corresponding to words which preserves hyponymy. Specifically, we introduce a composition rule called Compression, Compr, which is itself a completely positive map, and is therefore linear, generally non-commutative, and it preserves hyponymy. We give a number of proposals for the structure of Compr, based on spiders, cups, and caps, and generate a range of composition rules (see pages 16–20 of P10). As a particular case, we recover the previously introduced Mult, Fuzz and Phaser.

We test these rules on sentence entailment datasets from [KS16], and see some improvements over the performance of Fuzz and Phaser. We estimate the parameters of a simplified form of Compr based on entailment information from the aforementioned datasets, and find that while the learnt operator does not consistently outperform previously proposed mechanisms, it is competitive and has the potential to improve the use of a less simplified version.

2 Quantum magic squares

A *magic square* is a square matrix with positive entries such that every row and column sums to the same number (Fig. 1.6), and a *doubly stochastic matrix* is a magic square with real nonnegative entries where every row and column sums to 1. For example, dividing every entry of Dürer's magic square by 34 results in a doubly stochastic matrix. So doubly stochastic matrices contain a discrete probability measure in each row and each column.

It is well-known that the set of doubly stochastic matrices forms a polytope, whose vertices are the *permutation matrices*, i.e. doubly stochastic matrices with a single 1 in every row and column, and 0 elsewhere (that is, permutations of the identity matrix). This is the content of the famous Birkhoff–von Neumann Theorem.

A *quantum magic square* is a quantum generalization of a doubly stochastic matrix, where every number is promoted to a positive semidefinite matrix. So a quantum magic square contains a quantum measurement (i.e. a positive operator valued measurement, POVM) in every row and column (Definition 1 of P11). The normalisation conditions on the numbers (that they sum to 1) become the normalisation of the POVM (that they sum to the identity matrix).

A *quantum permutation matrix* is a quantum generalization of a permutation

Fig. 1.6: (Left) The magic square on the façade of the Sagrada Família in Barcelona, where every row and column adds to 33. (Right) The magic square in Albrecht Dürer's lithograph *Melencolia I*, where every row and column adds to 34.

matrix, where every 0 and 1 is promoted to an orthogonal projector (given that 0 and 1 are the only numbers that square to themselves). Thus, a quantum permutation matrix contains a projective valued measurement (PVM) in every row and column.

It is well-known that every POVM dilates to a PVM, by Naimark's Dilation Theorem. In words, every measurement can be 'purified', i.e. seen as part of a projective measurement in a larger system. But does this also hold for a two-dimensional array of POVMs? Namely,

Does every quantum magic square dilate to a quantum permutation matrix?

In other words: can every two-dimensional array of POVMs be dilated to a two-dimensional array of PVMs?

This is the main question addressed in P11, and the answer is 'no'—even in the simplest case where the internal positive semidefinite matrices are 2×2. This means that quantum magic squares do not admit the same kind of 'easy' characterisation as classical magic squares, that is, what happens at level 1 of the universe of Fig. 2 is deceivingly simple compared to what happens at higher levels. There must exist very strange (and thus very interesting) quantum magic squares. The proof relies on techniques from free semialgebraic geometry alluded to in page iv. A high level explanation of these results is provided in P13.

Ours is one contribution to the very lively subfield of quantum magic squares—recently an absolutely maximally entangled state of 4 systems with 6 levels each has been found [RBB$^+$21] (see this very nice Quanta magazine article about it). Absolutely maximally entangled states define a special kind of quantum magic square, and it is at present unclear how the results of [RBB$^+$21] relate to ours. Other recent very interesting developments include a quantum version of Sudoku [NP20].

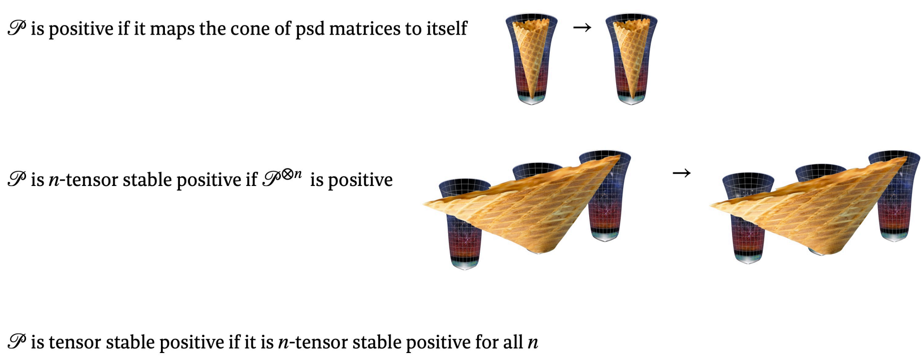

Fig. 1.7: A map $\mathscr{P}$ is positive if it maps to cone of positive semidefinite matrices to itself, n-tensor stable positive if $\mathscr{P}^{\otimes n}$ is positive, and tensor stable positive if n-tensor stable positive for all n. Note that $\mathscr{P}^{\otimes n}$ acts on n copies of the original universe, and $\mathscr{P}^{\otimes n}$ is positive if it preserves the cone on these n copies.

3 Shadows of infinity

A map $\mathscr{P}$ is *positive* if it maps the cone of positive semidefinite matrices to itself, and it is *tensor stable positive* if $\mathscr{P}^{\otimes n}$ is positive for all n [MHRW16] (Fig. 1.7). More formally, let $\mathscr{P} : \mathscr{M}_d \to \mathscr{M}_d$ be a linear map, where $\mathscr{M}_d$ is the space of $d \times d$ matrices. This map is positive if $\mathscr{P} : \mathrm{PSD}_d \to \mathrm{PSD}_d$, where PSD_d is the cone of positive semidefinite matrices of size $d \times d$. $\mathscr{P}^{\otimes n}$ is defined on the n-fold tensor product of $\mathscr{M}_d$,

$$\mathscr{P}^{\otimes n} : \underbrace{\mathscr{M}_d \otimes \ldots \otimes \mathscr{M}_d}_{n \text{ times}} \to \underbrace{\mathscr{M}_d \otimes \ldots \otimes \mathscr{M}_d}_{n \text{ times}},$$

and it is positive if it maps the cone of positive semidefinite matrices on the domain to itself,

$$\mathscr{P}^{\otimes n} : \mathrm{PSD}_{d^n} \to \mathrm{PSD}_{d^n}.$$

It is NP hard to decide whether a map is positive, and we conjecture that it is undecidable to decide if a map is tensor stable positive (P12, [vdEKS$^+$22]). We will revisit this conjecture in a couple of lines.

Now, completely positive maps, possibly followed by transposition, are tensor stable positive. These are the *trivial* tensor stable positive maps, whereas any non-trivial example is called *essential*. The central question considered in P12 is:

Are there any essential tensor stable positive maps?

This question was introduced and studied (but not settled) in [MHRW16], and we revisit it in P12 (but do not settle it either) by investigating it from two angles.

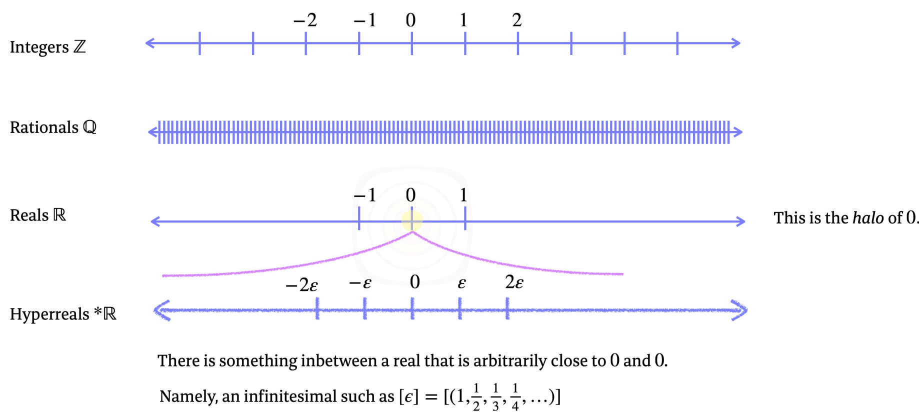

Fig. 1.8: The hyperreals are an extension of the reals, very much so as the reals are an extension of the rationals, and the rationals of the integers. Every real number, when considered as a hyperreal, contains a 'halo' of infinitesimally close elements. The figure shows the halo of the number 0. This halo is trivial in the reals—it is invisible, there is nothing in there—, but not in the hyperreals.

The first angle concerns undecidability. Consider the state

$$|\chi_n\rangle = \sum_{i_1,\ldots,i_n=1}^{d} |i_1, i_2\rangle \otimes |i_2, i_3\rangle \otimes \cdots \otimes |i_n, i_1\rangle, \tag{1.23}$$

where $|i\rangle$ denotes the i-th vector from the canonical orthonormal basis, and $|i_l, i_{l+1}\rangle$ is shorthand for $|i_l\rangle \otimes |i_{l+1}\rangle$. We show that, given $\mathscr{P}$, it is undecidable whether $\mathscr{P}^{\otimes n}(|\chi_n\rangle\langle\chi_n|)$ is positive semidefinite for all n. We conjecture that, given $\mathscr{P}$, it is undecidable whether $\mathscr{P}^{\otimes n}$ is a positive map for all n. We also show that if the conjecture were true there would exist essential tensor stable positive maps as well as bound entanglement with a negative partial transpose [MHRW16]. Note that here undecidability is a proof technique, instead of an end in itself.

The second angle concerns the hypercomplex numbers, for which we first need to define the hyperreals (Fig. 1.8). An integer is defined as an equivalence class of pairs of naturals numbers (which subtract to that integer), a rational as an equivalence class of pairs of integers (which give the same fraction), a real as an equivalence class of Cauchy sequences of rationals (which converge to that real), and a hyperreal as an equivalence class of sequences of real numbers, $^*\mathbb{R} = \mathbb{R}^{\mathbb{N}}/\sim$ [Gol98]. The equivalence $\sim$ is given by a so-called ultrafilter, which tells which elements of $\mathbb{R}^{\mathbb{N}}$ are equivalent to each other, and turns $^*\mathbb{R}$ into a field. In fact, $^*\mathbb{R}$ is a complete ordered field, meaning that it also has a total order, and that it contains its limits.

What is important for our purposes is that the reals are contained as a subfield

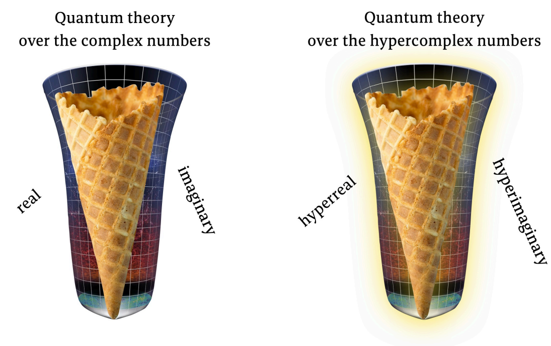

Fig. 1.9: We consider quantum theory over the hypercomplex numbers, where we describe quantum systems with matrices with hypercomplex entries (i.e. with a hyperreal and a hyperimaginary part). The halos provided by these hypercomplex numbers give rise to a 'glowing' version of usual quantum theory.

in the hyperreals: any real $r \in \mathbb{R}$ is embedded in the hyperreals as the equivalence class of the sequence $(r, r, r, r, \ldots)$. The situation is thus parallel to that of naturals being embedded in the integers, integers in the rationals, and rationals in the reals.

The hyperreals contain some fascinating objects[6], such as infinitesimals, which are smaller than than any positive real, yet larger than 0. For example, the equivalence class of the sequence $(1, 1/2, 1/3, 1/4, 1/5, \ldots)$ is such an infinitesimal. The infinitesimals around any hyperreal define its *halo*. The hyperreals also contain unlimited numbers, which are larger than any natural, yet not infinity. For example, the equivalence class of the sequence $(1, 2, 3, 4, 5, \ldots)$ is unlimited.

The hypercomplex $^*\mathbb{C}$ are the 'complexification' of the hyperreals, $^*\mathbb{C} =^* \mathbb{R}+i^*\mathbb{R}$, where i is the imaginary unit. They are an extension of the complex numbers, and they should not be confused with the also called hypercomplex numbers involving quaternions.

In P12 we consider quantum theory over the hypercomplex, instead of the complex (Fig. 1.9). This is a modification of quantum theory which, as far as we know, had not been considered before, and given our findings, should be considered from a foundational perspective, I believe.

In P12 we settle the main question regarding the existence of essential tensor stable positive maps in the positive, but for the hypercomplex numbers. Namely, we show that there exist essential tensor stable positive maps if the underlying field is that of the hypercomplex, instead of the complex. From here it follows that there exist undistillable quantum states with a negative partial transpose (NPT) over the hypercomplex. In other words, the existence of NPT bound (i.e. undistillable) entanglement is an open problem in quantum information (see e.g. [HRŻ20]), which

[6]The reals, the rationals, the integers and the naturals do too!

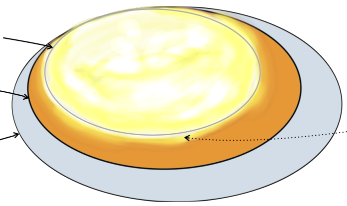

Fig. 1.10: There exist undistillable NPT quantum states in the hypercomplex. This follows from the fact that, in the hypercomplex, states with a positive partial transpose (PPT) glow so that they become NPT but still remain undistillable (P12). 'Glowing' means that they have these undistillable NPT states in their halo (see Fig. 1.9). This halo vanishes in the complex, so we cannot conclude anything about the usual complex case.

can be solved over the hypercomplex, as P12 shows. These statements follow from the fact that the halo of trivial tensor stable positive maps contains essential tensor stable positive maps (Fig. 1.10). The halo becomes trivial when the hypercomplex are replaced by the complex, so unfortunately our result does not allow to conclude anything about the 'usual' situation.

Bibliography

[AKLT87] I. Affleck, T. Kennedy, E. H. Lieb, and H. Tasaki. Rigorous results on valence-bond ground states in antiferromagnets. *Phys. Rev. Lett.*, 59:799, 1987.

[BCLM18] D. Bankova, B. Coecke, M. Lewis, and D. Marsden. Graded hyponymy for compositional distributional semantics. *JLM*, 6:225, 2018.

[BSM03] A. Berman and N. Shaked-Monderer. *Completely Positive Matrices*. World Scientific, 2003.

[CGLM08] P. Comon, G. Golub, L.-H. Lim, and B. Mourrain. Symmetric tensors and symmetric tensor rank. *SIAM J. Matrix Anal. Appl.*, 30:1254, 2008.

[CGW11] X. Chen, Z.-C. Gu, and X.-G. Wen. Classification of gapped symmetric phases in one-dimensional spin systems. *Phys. Rev. B*, 83:035107, 2011.

[CM20] B. Coecke and K. Meichanetzidis. Meaning updating of density matrices. *arXiv:2001.00862*, 2020.

[CPGSV17] J. I. Cirac, D. Perez-Garcia, N. Schuch, and F. Verstraete. Matrix Product Density Operators: Renormalization Fixed Points and Boundary Theories. *Ann. Phys.*, 378:100, 2017.

[CPGSV21] I. Cirac, D. Perez-Garcia, N. Schuch, and F. Verstraete. Matrix Product States and Projected Entangled Pair States: Concepts, Symmetries, and Theorems. *Rev. Mod. Phys.*, 93:045003, 2021.

[CR93] J. E. Cohen and U. G. Rothblum. Nonnegative ranks, decompositions, and factorizations of nonnegative matrices. *Linear Algebra Appl.*, 190:149, 1993.

[CSC10] B. Coecke, M. Sadrzadeh, and S. Clark. Mathematical foundations for a compositional distributional model of meaning. 2010.

[CVZ19] M. Christandl, P. Vrana, and J. Zuiddam. Asymptotic tensor rank of graph tensors: beyond matrix multiplication. *Computational complexity*, 28:57, 2019.

[DCC+16] G. De las Cuevas, T. S. Cubitt, J. I. Cirac, M. M. Wolf, and D. Perez-Garcia. Fundamental limitations in the purifications of tensor networks. *J. Math. Phys.*, 57:071902, 2016.

[Den89] Denisov, L. V. Infinitely divisible Markov mappings in quantum probability theory. *Theory Probab. Appl.*, 33:392, 1989.

[DKN22] G. De las Cuevas, A. Klingler, and T. Netzer. Border ranks of general tensor decompositions with invariance and positivity. *In preparation0*, 2022.

[DSPGC13] G. De las Cuevas, N. Schuch, D. Perez-Garcia, and J. I. Cirac. Purifications of multipartite states: limitations and constructive methods. *New J. Phys.*, 15:123021, 2013.

[FGP+15] H. Fawzi, J. Gouveia, P. A. Parrilo, R. Z. Robinson, and R. R. Thomas. Positive semidefinite rank. *Math. Program.*, 15:133, 2015.

[FMP+12] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf. Linear vs. Semidefinite Extended Formulations: Exponential Separation and Strong Lower Bounds. *STOC '12 Proc. of the 44th Symposium on Theory of Computing*, page 95, 2012.

[FNW92] M. Fannes, B. Nachtergaele, and R. F. Werner. Finitely correlated states on quantum spin chains. *Commun. Math. Phys.*, 144:443, 1992.

[Gol98] R. Goldblatt. *Lectures on the hyperreals*. Springer, 1998.

[GPT13] J. Gouveia, P. A. Parrilo, and R. R. Thomas. Lifts of convex sets and cone factorizations. *Math. Oper. Res.*, 38:248, 2013.

[HCOV13] J. Haegeman, J. I. Cirac, T. J. Osborne, and F. Verstraete. Calculus of continuous matrix product states. *Phys. Rev. B*, 88:85118, 2013.

[HRŻ20] P. Horodecki, Ł. Rudnicki, and K. Życzkowski. Five open problems in quantum information. *arXiv:2002.03233*, 2020.

[Iva19] G. M. Ivanov. Approximate Caratheodory's theorem in uniformly smooth Banach spaces. 2019.

[KGE14] M. Kliesch, D. Gross, and J. Eisert. Matrix product operators and states: NP-hardness and undecidability. *Phys. Rev. Lett.*, 113:160503, 2014.

[Kho87] A. S. Kholevo. Infinitely divisible measurements in quantum probability theory. *Theory Probab. Appl.*, 31:493, 1987.

[Koz97] D. C. Kozen. *Automata and Computability*. Springer, 1997.

[KS16] D. Kartsaklis and M. Sadrzadeh. Distributional Inclusion Hypothesis for Tensor-based Composition. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 2849–2860, 2016.

[Lew19] M. A. F. Lewis. Compositional hyponymy with positive operators. *Proc. recent Advances in NLP*, 2019.

[LQY11] J. M. Landsberg, Y. Qi, and K. Ye. On the geometry of tensor network states. *arXiv: 1105.4449*, 2011.

[MHRW16] A. Müller-Hermes, D. Reeb, and M. M. Wolf. Positivity of linear maps under tensor powers. *J. Math. Phys.*, 57:015202, 2016.

[NP20] I. Nechita and J. Pillet. SudoQ – a quantum variant of the popular game. *arXiv:2005.10862*, 2020.

[Oru19] R. Orus. Tensor networks for complex systems. *Nat. Rev. Phys.*, 2019.

[PGVWC07] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac. Matrix Product State representations. *Quantum Inf. Comput.*, 7:401, 2007.

[PKCS15] R. Piedeleu, D. Kartsaklis, B. Coecke, and M. Sadrzadeh. Open System Categorical Quantum Semantics in Natural Language Processing. In L. S. Moss and P. Sobocinski, editors, *6th Conference on Algebra and Coalgebra in Computer Science, Nijmegen, The Netherlands*, volume 35 of *LIPIcs*, pages 270–289. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2015.

[RB01] R. Raussendorf and H. J. Briegel. A One-Way Quantum Computer. *Phys. Rev. Lett.*, 86:5188, 2001.

[RBB+21] S. A. Rather, A. Burchardt, W. Bruzda, G. Rajchel-Mieldzioć, A. Lakshminarayan, and K.Życzkowski. Thirty-six entangled officers of Euler: Quantum solution to a classically impossible problem. *arXiv:2104.05122*, 2021.

[SDD20] S. Stengele, D. Drexel, and G. De las Cuevas. Classical spin Hamiltonians are context-sensitive languages. *arXiv:2006.03529*, 2020.

[Shi18] Y. Shitov. A counterexample to Comon's conjecture. *SIAM J. Appl. Alg. Geometry*, 2(3):428–443, 2018.

[SPGC11] N. Schuch, D. Perez-Garcia, and I. Cirac. Classifying quantum phases using matrix product states and projected entangled pair states. *Phys. Rev. B*, 84:165139, 2011.

[VC10] F. Verstraete and J. I. Cirac. Continuous Matrix Product States for Quantum Fields. *Phys. Rev. Lett.*, 104:190405, 2010.

[vdEKS+22] M. van der Eyden, A. Klingler, S. Stengele, T. Reinhart, and G. De las Cuevas. From undecidability to NP-completeness and back. *In preparation*, 2022.

[Vid03] G. Vidal. Efficient classical simulation of slightly entangled quantum computation. *Phys. Rev. Lett.*, 91:147902, 2003.

[Wol11] M. M. Wolf. Quantum Channels and Operations: Guided Tour. *Unpublished lecture notes*, 2011.