Internationale Mathematische Nachrichten
International Mathematical News
Nouvelles Mathématiques Internationales

Die IMN wurden 1947 von R. Inzinger als „Nachrichten der Mathematischen Gesellschaft in Wien“ gegründet. 1952 wurde die Zeitschrift in „Internationale Mathematische Nachrichten“ umbenannt und war bis 1971 offizielles Publikationsorgan der „Internationalen Mathematischen Union“.

Von 1953 bis 1977 betreute W. Wunderlich, der bereits seit der Gründung als Redakteur mitwirkte, als Herausgeber die IMN. Die weiteren Herausgeber waren H. Vogler (1978–79), U. Dieter (1980–81, 1984–85), L. Reich (1982–83), P. Flor (1986–99), M. Drmota (2000–2007) und J. Wallner (2008–2017).

Herausgeber:

Österreichische Mathematische Gesellschaft, Wiedner Hauptstraße 8–10/104, A-1040 Wien. email *imn@oemg.ac.at*, *http://www.oemg.ac.at/*

Redaktion:

C. Fuchs (Univ. Salzburg, Herausgeber)
H. Humenberger (Univ. Wien)
R. Tichy (TU Graz)
J. Wallner (TU Graz)

Bezug:

Die IMN erscheinen dreimal jährlich und werden von den Mitgliedern der Österreichischen Mathematischen Gesellschaft bezogen.

Jahresbeitrag: € 35,–

Bankverbindung:
IBAN AT83-1200-0229-1038-9200 bei der Bank Austria-Creditanstalt (BIC-Code BKAUATWW).

Eigentümer, Herausgeber und Verleger: Österr. Math. Gesellschaft. Satz: Österr. Math. Gesellschaft. Druck: Weinitzendruck, 8044 Weinitzen.

ISSN 0020-7926

Österreichische Mathematische Gesellschaft

Gegründet 1903
http://www.oemg.ac.at/
email: *oemg@oemg.ac.at*

Sekretariat:

Alpen-Adria-Universität Klagenfurt,
Institut für Mathematik
Universitätsstraße 65-67
A-9020 Klagenfurt
email: *oemg@oemg.ac.at*

Vorstand des Vereinsjahres 2019:

B. Kaltenbacher (Univ. Klagenfurt):
Vorsitzende
J. Wallner (TU Graz):
Stellvertretender Vorsitzender
C. Fuchs (Univ. Salzburg):
Herausgeber der IMN
M. Ludwig (TU Wien):
Schriftführerin
M. Haltmeier (Univ. Innsbruck):
Stellvertretender Schriftführer
B. Lamel (Univ. Wien):
Kassier
P. Grohs (Univ. Wien):
Stellvertretender Kassier
E. Buckwar (Univ. Linz):
Beauftragte für Frauenförderung
C. Heuberger (Univ. Klagenfurt):
Beauftragter f. Öffentlichkeitsarbeit

Beirat:

A. Binder (Linz)
M. Drmota (TU Wien)
H. Edelsbrunner (ISTA)
H. Engl (Univ. Wien)
H. Heugl (Wien)
W. Imrich (MU Leoben)
M. Kim (MathWorks)
M. Koth (Univ. Wien)
C. Krattenthaler (Univ. Wien)
M. Oberguggenberger (Univ. Innsbruck)
W. Müller (Univ. Klagenfurt)
H. Niederreiter (ÖAW)
W. G. Nowak (Univ. Bodenkultur)
W. Schachermayer (Univ. Wien)
K. Sigmund (Univ. Wien)
H. Sorger (Wien)
R. Tichy (TU Graz)
H. Zeiler (Wien)

Vorsitzende von Sektionen und Kommissionen:

W. Woess (Graz)
H.-P. Schröcker (Innsbruck)
C. Pötzsche (Klagenfurt)
F. Pillichshammer (Linz)
S. Blatt (Salzburg)
I. Fischer (Wien)
H. Humenberger (Didaktikkommission)
W. Müller (Verantwortlicher für Entwicklungszusammenarbeit)

Die Landesvorsitzenden und der Vorsitzende der Didaktikkommission gehören statutengemäß dem Beirat an.

Mitgliedsbeitrag:

Jahresbeitrag: € 35,–

Bankverbindung: IBAN AT83-1200-0229-1038-9200

Internationale Mathematische Nachrichten

International Mathematical News

Nouvelles Mathématiques Internationales

Nr. 240 (73. Jahrgang) **April 2019**

Inhalt

Die Formel auf der Titelseite entstammt einem Überblicksartikel von Peter Scholze über seine neuen Ergebnisse aus der p-adischen Geometrie anlässlich seines ICM-Vortrags 2018 in Rio de Janeiro (siehe https://arxiv.org/pdf/1712.03708.pdf). Die Formel beantwortet eine Frage von Tate, nämlich, dass eine eindeutige Galois-äquivariante Zerlegung der aufliegenden Filtrierung existiert, die zu einem in der Formel angegebenen Galois-äquivariantem Isomorphismus führt. Dabei ist X der Basiswechsel eines eigentlichen glatten rigiden Raums, definiert über einem diskret bewerteten Raum $K \subset C$. Der angegebene Isomorphismus existiert jedoch nicht in Familien analog zur Hodge-Zerlegung über den komplexen Zahlen, die ebenfalls nicht holomorph in Familien variiert.

Internat. Math. Nachrichten
Nr. 240 (2019), 1–10

Quantenlogiken und Numerische Ereignisse

Dietmar Dorninger
TU Wien

1 Einleitung

Quantenlogik ist kein wohldefinierter Begriff. Quantenlogiken werden auf unterschiedliche Weise im mathematisch-physikalischen Kontext beschrieben, sie werden in der Logik als dreiwertige oder temporäre Logiken behandelt und auch vom Standpunkt der Linguistik und Philosophie aus betrachtet – jeweils auf verschiedenen konzeptionellen Auffassungen aufbauend. Gemeinsam ist allen Begriffserklärungen, dass den Anstoß dazu Phänomene der Quantenmechanik geben, welche im Widerspruch zu unseren im Alltag als logisch empfundenen Erfahrungen stehen.

Wir gehen hier von dem von John von Neumann und Garret Birkhoff 1936 geprägten Begriff „Quantenlogik“ aus (siehe [6]), der für unzählige weitere Untersuchungen Pate gestanden hat, und unter den zahlreichen Fragestellungen, denen dann nachgegangen worden ist, greifen wir das Problem heraus, wie man aufgrund von Messergebnissen feststellen kann, ob es sich bei einem physikalischen System um ein quantenmechanisches System handelt oder eines, das der klassischen Mechanik zuzurechnen ist. Von den vielen Möglichkeiten, die es wiederum hierzu gibt, ist wohl das Studium von Bellschen Ungleichungen, die 1964 von J. S. Bell angegeben und später von anderen Autoren verallgemeinert worden sind (siehe [1]), am bekanntesten. Wir wollen hier aber einen anderen Weg aufzeigen, nämlich, wie man von Messergebnissen direkt Rückschlüsse auf die zugrunde liegende Quantenlogik im Sinne von J.von Neumann und G. Birkhoff ziehen kann, ohne von vornherein nähere Annahmen über die Quantenlogik zu machen, was bei quantenmechanischen Berechnungen oft so nicht der Fall ist. Hierzu stellen wir den vom Physiker E. Beltrametti gemeinsam mit dem Mathematiker M.

ISSN 0020-7926

Mączyński 1979 eingeführten Begriff „Numerisches Ereignis“ vor (siehe [3],[4]). Numerische Ereignisse sind im Wesentlichen n-tupel von Wahrscheinlichkeiten von Messergebnissen, welche partiell geordnet Strukturen ergeben, für die man Kriterien finden kann, ob man es mit einer klassischen oder quantenmechanischen Problemstellung zu tun hat. Wie bedeutsam derartige Fragestellungen sind, kann man bereits an dem einfachen Beispiel ersehen, dass es bei der fortschreitenden Miniaturisierung von Computerchips passieren kann, dass 0 und 1 nicht mehr korrekt ausgelesen werden; sie könnten sich bereits in einem für die Quantenmechanik charakteristischen überlagerten Zustand befinden.

2 Quantenlogiken im Sinne von J. von Neumann und G. Birkhoff

Wesentlich zum Verständnis dieser Definition von Quantenlogiken und all ihren Verallgemeinerungen ist der algebraische Begriff einer orthomodularen Halbordnung. Darunter ist eine halbgeordnete Menge $(L,\leq)$ zu verstehen, welche ein kleinstes Element 0 und ein größtes Element 1 hat, in der es eine einstellige Operation $'$ gibt und für die mit den Bezeichnungen $a\wedge b$ und $a\vee b$ für die größte untere bzw. kleinste obere Schranke zweier Elemente a und b in L gilt:

(O1) $a\wedge a'=0$, $a\vee a'=1$, $a''=a$ for all $a\in L$,

(O2) aus $a\leq b$ folgt $b'\leq a'$ for all $a,b\in L$,

(O3) falls $a\leq b'$, in welchem Fall a und b orthogonal heißen – in Zeichen $a\perp b$, dann existiert $a\vee b$ in L,

(O4) für $a\leq b$, $a,b\in L$, gilt $b=a\vee(a'\wedge b)$ (orthomodulares Gesetz).

Falls in einer orthomodularen Halbordnung L für alle $a,b\in L$ sowohl das Infimum $a\wedge b$ als auch das Supremum $a\vee b$ existiert, so heißt L (wenn man die Infimums- und Supremumsbildung als zweistellige Operationen auffasst) ein orthomodularer Verband. Gelten überdies die Distributivgesetze $a\wedge(b\vee c)=(a\wedge b)\vee(a\wedge c)$ und $a\vee(b\wedge c)=(a\vee b)\wedge(a\vee c)$ für alle $a,b,c\in L$, so handelt es sich um eine Boolesche Algebra.

Das für uns hier wichtigste Beispiel für einen orthomodularen Verband ist der Verband der abgeschlossenen Unterräume eines separablen Hilbertraums H bzw. der dazu isomorphe Verband der Projektoren von H, welchen wir mit $P(H)$ bezeichen wollen.

In ihrer richtungsweisenden Arbeit 1936 haben von Neumann und Birkhoff jedem quantenmechanischen System einen Hilbertraum H zugeordnet, wobei dessen Elemente ψ mit Norm $\|\psi\|=1$ den sog. reinen Zuständen des Systems entsprechen und die Observablen (beobachtbaren Größen) $\mathcal{A}$ durch selbstadjungierte

lineare Operatoren A repräsentiert werden. Schreiben wir $\langle \varphi, \psi \rangle$ für das innere Produkt von $\varphi, \psi \in H$, so stellt bei diesen Entsprechungen $\langle A\psi, \psi \rangle$ den Erwartungswert für den Messwert der Observablen $\mathcal{A}$ dar, wenn sich das physikalische System im Zustand ψ befindet, und die Eigenwerte zu den Eigenzuständen von A geben die möglichen exakten Werte von $\mathcal{A}$ wieder. Ordnet man einem Operator A gemäß dem Spektralsatz sein Spektralmaß α zu, also eine Abbildung von den Borelmengen $\mathcal{B}(\mathbb{R})$ in den Verband $P(H)$ der Projektoren von H, so kann $\langle \alpha(E)(\psi), \psi \rangle$ im quantenmechanischen Modell als die Wahrscheinlichkeit dafür interpretiert werden, dass eine Messung von $\mathcal{A}$ im Zustand ψ zu einem innerhalb von $E \in \mathcal{B}(\mathbb{R})$ liegenden Wert führt.

Jeder Projektor P_U auf einen (nicht-trivialen) abgeschlossenen Unterraum U von H ist seinerseits ein selbstadjungierter linearer Operator von H und hat als Eigenwerte nur 0 und 1, d.h. die durch $P \in P(H)$ dargestellte Observable $\mathcal{P}$ kann als eine „Frage“ an das quantenmechanische System mit den möglichen Antworten „ja“ für den Eigenwert „1“ und „nein“ für den Eigenwert 0 gedeutet werden. Ferner kann die Observable $\mathcal{P}'$, welche durch den Projektor auf den zu U orthogonalen Unterraum definiert ist, als Negation der Frage $\mathcal{P}$ angesehen werden. Sind $\mathcal{P}$ und $\mathcal{Q}$ Fragen und folgt aus der Annahme, dass $\mathcal{P}$ den Wert 1 hat, dass auch $\mathcal{Q}$ den Wert 1 hat, so kann man im Sinne einer Art logischen Kalküls sagen $\mathcal{P}$ impliziert $\mathcal{Q}$. Das erklärt, warum $P(H)$ als Quantenlogik bezeichnet wurde. Ist H unendlichdimensional, so spricht man von der Standardlogik der Quantenmechanik. Weil von Neumann und Birkhoff Fragen auch als experimentelle Aussagen über Ereignisse angesprochen haben, hat sich für $P(H)$ und dessen Verallgemeinerungen neben Quantenlogik auch der Begriff „Ereignisraum“ eingebürgert.

Zur Illustration letzterer Begriffsbildung betrachten wir (im Sinn eines sog. Stern-Gerlach-Experiments) den Spin von Elektronen in einem Elektronenstrahl, welcher sich durch einen Strahlteiler in einen Strahl von Elektronen mit Spin „nach oben“ und einen solchen mit Spin „nach unten“ aufspalten lässt, was auf einem Beobachtungsschirm verifiziert werden kann. Interpretieren wir den Spin als Vektor im $\mathbb{R}^2$, welchen wir in einem kartesischen x, y-Koordinatensystem auf die x-und y-Achse projizieren und dessen Komponenten wir dann ebenfalls die Eigenschaften Spin „nach oben“ und Spin „nach unten“ zuordnen, dann zeigt sich das folgende Phänomen: Schalten wir drei Strahlteiler hintereinander, wobei der erste die x-Komponente „nach unten“ völlig ausfiltert, senden den so um die x-Komponente „nach unten“ reduzierten Strahl in den zweiten Strahlteiler, der die y-Komponente „nach oben“ ausfiltert und lassen anschließend den Strahl einen dritten Strahlteiler passieren, der die x-Komponente in Spin „nach oben“ und „nach unten“ teilt, so erwartet man, dass auf einem hinter dem dritten Strahlteiler liegenden Beobachtungsschirm nur noch Elektronen mit Spin „nach oben“ vorkommen können. Tatsächlich aber können in gleichem Maße wieder Elektronen mit der x-Komponente Spin „nach unten“ wie „nach oben“ nachgewiesen werden. Bezeichnet A das Ereignis, dass die x-Komponente einen Spin „nach oben“ aufweist und

B das Ereignis, dass die y-Komponente Spin „nach oben“ hat, und schreiben wir $'$ für die Komplementbildung, $\wedge$ für „und“ sowie $\vee$ für „oder“, so kann man den Ausgang unseres Experiments als $A \wedge B' = (A \wedge B') \vee (A' \wedge B')$ interpretieren. Für die vorkommenden Ereignisse ist dann das Distributivgesetz nicht erfüllt, welches $A \wedge B' = B'$ zur Folge hätte. Als Ereignisraum kann also keine Boolesche Algebra infrage kommen; ein geeigneter Ereignisraum ist in Abb.1 wiedergegeben.

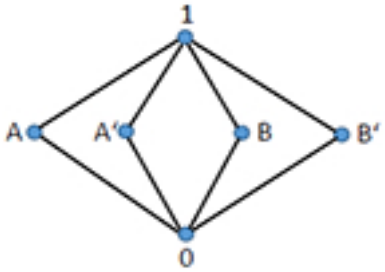

Abbildung 1: Quantenlogik zum Stern-Gerlach-Experiment

Kehren wir nochmals zu unserer Hilbertraumlogik $P(H)$ zurück. Für einen festen Vektor $\psi \in H$ mit Norm $\|\psi\| = 1$, d.h. also für einen reinen Zustand ψ, und einen Projektor $P \in P(H)$ ist $\mu_\psi(P) := \langle P\psi, \psi \rangle$ die Wahrscheinlichkeit dafür, dass $\mathcal{P}$ im Zustand ψ den Wert 1 hat. Davon verallgemeinert man: Eine beliebige Abbildung $\mu : P(H) \to [0,1]$ mit $\mu(H) = 1$ und $\mu(\sum_i P_i) = \sum_i \mu(P_i)$ für eine Folge paarweise orthogonaler Projektoren P_i, $i = 1, 2, \ldots$, heiße ein Wahrscheinlichkeitsmaß auf $P(H)$. Dies aus gutem Grund, denn ein charakteristisches Merkmal der klassischen Mechanik ist, dass zwei beliebige Observable kommutieren, was im quantenmechanischen Modell bedeutet, dass $P(H)$ eine Boolesche Algebra (σ-Algebra) ist und μ dann eine klassische Wahrscheinlichkeitsverteilung darstellt.

Sehr bald schon wurde das Hilbertraum-Modell der Quantenmechanik verallgemeinert. Heute wird unter Quantenlogik oft eine beliebige orthomodulare Halbordnung L verstanden, welche, wie erwähnt, auch als Ereignisraum angesprochen wird. Eine Abbildung μ von L in $[0,1]$ mit der Eigenschaft, dass $\mu(1) = 1$ und $\mu(a \vee b) = \mu(a) + \mu(b)$ für $a \perp b$, $a, b \in L$ wird in Anspielung an die Zustände ψ im Hilbertraummodell ein Zustand genannt, und μ heißt ein Wahrscheinlichkeitsmaß auf L, falls $\mu(1) = 1$ und $\mu(\sum_i a_i) = \sum_i \mu(a_i)$ für jede Folge von paarweise orthogonalen Elementen $a_i \in L$. Eine wichtige Rolle spielt bei allen Verallgemeinerungen die Existenz von Mengen von Zuständen, welche die Halbordnungsrelation $\leq$ von L festlegen und ordnungsbestimmend oder voll heißen. Dabei wird eine Menge S von Zuständen als voll bezeichnet, falls $a \leq b$ in L genau dann gilt, wenn $\mu(a) \leq \mu(b)$ für alle $\mu \in S$. Auch diese Definition geht auf eine Eigenschaft der Hilbertraum-Logik $P(H)$ zurück.

3 Algebren von numerischen Ereignissen

Sei S eine Menge von Zuständen eines physikalischen Systems und $p(s)$ die Wahrscheinlichkeit, dass ein gewisses Ereignis eintritt, wenn sich das System im Zustand $s \in S$ befindet. Betrachtet man $p(s)$ für alle $s \in S$, so erhält man eine Funktion $p : S \to [0,1]$, welche E. Beltrametti und M. Mączyński 1991 ein numerisches Ereignis oder eine multidimensionale Wahrscheinlichkeit bzw. zur Präzisierung, S-Wahrscheinlichkeit, genannt haben. Später wurde darauf verwiesen, dass ein numerisches Ereignis $p(s)$ auch als ein Spezialfall der von G. Mackey im Rahmen seiner axiomatischen Einführung in die Grundlagen der Quantenmechanik verwendeten Funktion $p(\mathcal{A}, s, E)$ angesehen werden kann, wenn man $\mathcal{A}$ und E festhält. $p(\mathcal{A}, s, E)$ ist die Wahrscheinlichkeit dafür, dass eine Messung von $\mathcal{A}$ beim Zustand s zu einem Wert in der Borelmenge E führt (siehe [14]).

Gegeben sei nun eine Menge W von S-Wahrscheinlichkeiten, welche die konstanten Funktionen 0 und 1 enthält. Um die Menge W in einen Zusammenhamng mit Quantenlogiken zu bringen, gingen Beltrametti und Mączyński von einem unbekannten Ereignisraum L aus, den sie als orthomodulare Halbordnung annahmen, und definierten ein S-Wahrscheinlichkeitsmaß auf L als eine Abbildung m von L in W mit den Eigenschaften

(i) $m(0) = 0, m(1) = 1$

(ii) $m(a') = 1 - m(a)$

(iii) $m(a_1 \vee a_2) = m(a_1) + m(a_2)$ für $a_1, a_2 \in L$ mit $a_1 \perp a_2$.

Dabei sind 0 und 1 als Argument von m das kleinste bzw. größte Element von L, und $+$ und $-$ sind die Addition und Subtraktion in $\mathbb{R}$. Gilt für alle $a, b \in L$, dass $m(a) \leq m(b)$ äquivalent zu $a \leq b$ ist, so heißt m vollständig.

Schreiben wir für festes $s \in S$ $m_s(a)$ statt $m(a)(s)$ und erweitern (iii) zur σ-Additivität, so sieht man unmittelbar, dass die Abbildung $a \to m_s(a)$ ein Wahrscheinlichkeitsmaß auf L ist. (Über Wahrscheinlichkeitsmaße auf Quantenlogiken bzw. Ereignisräumen gibt es eine weitläufige Literatur, von der wir die in direktem Zusammenhang mit diesen Ausführungen stehenden Bücher [2],[14] hervorheben.). Wesentlich für das Weitere ist nun der

Satz 3.1. *(siehe [3]) Eine Menge W von S-Wahrscheinlichkeiten ist das Bild eines vollständigen S-Wahrscheinlichkeitsmaßes genau dann, wenn W die folgenden Eigenschaften hat:*

(A1) $0 \in W$

(A2) $p \in W \Longrightarrow 1 - p \in W$

(A3) *Sind $p,q,r \in W$ paarweise orthogonal, so ist $p+q+r \in W$.*

Eine Menge W von S-Wahrscheinlichkeiten, welche (A1)–(A3) erfüllt, wurde zunächst ([4], 1993) „Raum numerischer Ereignisse“ und später ([5], 2007) „Algebra von S-Wahrscheinlichkeiten“ bzw. „Algebra von numerischen Ereignissen“ genannt (wobei es sich genauer eigentlich um partielle Algebren handelt). Das hervorstechenste Merkmal dieser Algebren von numerischen Ereignissen ist, dass man im Wesentlichen mit den Rechenoperationen $+$ und $-$ und der Relation $\leq$ auskommt, um Eigenschaften zu charakterisieren und Rechenverfahren zu formulieren. Um welche Art von Halbordnungen es sich bei den Algebren von numerischen Ereignissen vom algebraischen Standpunkt aus handelt, folgt aus einem bereits 1973 publizierten Satz von M. Mączyński and T. Traczyk (siehe [15]), welcher in unserer gegenwärtigen Notation lautet:

Satz 3.2. *Eine Algebra von numerischen Ereignissen ist bezüglich der partiellen Ordnung von Funktionen eine orthomodulare Halbordnung, für die es eine volle Menge von Zuständen gibt, und umgekehrt ist jede solche Halbordnung isomorph zu einer Algebra von numerischen Ereignissen.*

Aufgrund der Sätze 3.1 und 3.2 sind Algebren von numerischen Ereignissen epimorphe Bilder von Quantenlogiken mit einer vollen Menge von Zuständen und eignen sich daher selbst als Quantenlogiken. Unter diesen sind auch alle Hilbertraum-Quantenlogiken zu finden, wenn man den folgenden Zusammenhang herstellt: Sei $S(H)$ die Menge aller eindimensionalen Unterräume des Verbandes $P(H)$ der abgeschlossenen Unterrräume eines separablen Hilbertraums H und $p : S(H) \to [0,1]$, definiert durch $p(s) := \langle P\psi_s, \psi_s \rangle$, wobei ψ_s ein beliebiger Einheitsvektor des Unterraums $s \in S(H)$ ist. Dann bildet die Menge aller so definierten S-Wahrscheinlichkeiten p eine Algebra von numerischen Ereignissen und, wie man zeigen kann, ist diese zu $P(H)$ isomorph. Die Definition von S-Wahrscheinlichkeitsmaßen hatte auch dazu geführt, im Umkehrschluss beliebige orthomodulare Halbordnungen als Ereignisräume (bzw. Quantenlogiken) zu betrachten und dann Observable als S-Wahrscheinlichkeitsmaße auf diesen zu interpretieren (vgl. [13]).

Als Nächstes wollen wir zwei Fragen klären: Erstens, wie kann man von einer Algebra von S-Wahrscheinlichkeiten feststellen, dass diese eine Boolesche Algebra ist, woraus man dann schließen kann, dass es sich um ein klassisches System handelt, und zweitens, wenn nur eine gewisse Anzahl von Messergebnissen vorliegt, welche keinen Schluss auf die gesamte Algebra von numerischen Ereignissen zulässt, wie kann man dann vorgehen? Lässt sich die gefundene Menge von S-Wahrscheinlichkeiten in eine Boolesche Algebra einbetten oder ist dies unmöglich?

4 Quantenmechanik oder klassische Mechanik

Angenommen, es liegt eine Algebra von S-Wahrscheinlichkeiten W vor. Dann fragen wir uns natürlich zunächst, woher wir ganz W kennen, auch wenn keine oder nur wenige Messdaten zur Verfügung stehen. Dabei kann die Kenntnis von W auf zweierlei Wegen zustandegekommen sein:

Erstens, aufgrund von bekannten physikalischen Gesetzen oder Modellvorstellungen, wie etwa beim Beispiel der Bestimmung der Wahrscheinlichkeit $p(s)$, dass ein Photon einen Polarisator im Zustand s durchquert: Seien s_1 und s_2 die Zustände einer linearen Polarisation entlang der x-bzw. y-Achse eines kartesischen (x,y,z)-Koordinatensystems, in dessen z-Richtung sich das Photon bewegt. Verdreht man den Polarisator um den Winkel γ, so erhält man nach dem (physikalischen) Gesetz von Malus mit $\delta = cos^2(\gamma)$ im Hinblick auf die x- und y-Achse insgesamt die folgende Menge von $\{s_1,s_2\}$-Wahrscheinlichkeiten: $W = \left\{(0,0),(1,0),(0,1),(\delta,1-\delta),(1-\delta,\delta),(1,1)|0<\delta<\frac{1}{2}\right\}$, welche eine Algebra von $\{s_1,s_2\}$-Wahrscheinlichkeiten ist. (Diese ist isomorph zum Verband der Unterräume des Euklidischen Raumes $\mathbb{R}^2$.)

Eine weitere Möglichkeit, um Kenntnis einer Algebra von S-Wahrscheinlichkeiten zu gewinnen, besteht darin, dass man, ausgehend von einer Menge V von S-Wahrscheinlichkeiten, welche man durch Messungen erhalten hat, auf algebraischem Weg die kleinste Algebra von S-Wahrscheinlichkeiten W bestimmt, welche V als Unterstruktur enthält.

Liegt nun eine Algebra von S-Wahrscheinlichkeiten vor, gilt es, Kriterien dafür zu finden, dass diese eine Boolesche Algebra ist.

Für $|S| = 2$ ist dies sehr einfach, denn dann sind alle Algebren von S-Wahrscheinlichkeiten Verbände, die so wie bei dem obigen Beispiel vom Polarisator aus einer Kette von paarweise unvergleichbaren Elementen zwischen 0 und 1 bestehen, und die einzigen Boolesche Algebren darunter sind diejenigen mit zwei und vier Elementen ([7]).
Für $|S| > 2$ gibt es zahlreiche Charakterisierungen, von denen wir eine anführen, die zu einem einfachen Rechenverfahren führt, das dann, wenn das Kriterium nicht erfüllt ist, rasch abbrechen kann.

Satz 4.1. *([9]) Eine Algebra von S-Wahrscheinlichkeiten W ist genau dann eine Boolesche Algebra, wenn es zu allen $p,q \in W$ zwei (eindeutige) orthogonale Elemente r,s gibt mit $p \perp r$, $q \perp r$ und $p+r = s+q$.*

Da eine Boolesche Algebra insbesondere ein orthomodularer Verband ist, ist es auch von Interesse, verbandsgeordnete Algebren von numerischen Ereignissen als Boolesche Algebren zu charakterisieren ([8]). Ferner lassen sich auch für den Fall, dass eine Algebra von S-Wahrscheinlichkeiten eine mengentheoretische Darstellung besitzt, d.h., dass eine sog. konkrete Quantenlogik (siehe [16]) vorliegt,

Kriterien dafür angeben (siehe etwa [12]), dass es sich um eine Boolesche Algebra handelt. Im Hinblick auf ableitbare Rechenverfahren eignen sich die meisten Sätze hierbei aber in erster Linie für den Nachweis, dass eine Bedingung nicht erfüllt ist, in welchem Fall es dann naheliegt, anzunehmen, dass man es mit keinem klassischen physikalischen System zu tun hat.

Liegt nur eine (kleine) Menge von numerischen Ereignissen vor, so möchte man aus dieser unmittelbar herauslesen können, ob es sich um Messwerte handelt, die der klassischen Mechanik zuzurechnen sind oder nicht. In diesem Fall kann man für S-Wahrscheinlichkeiten, die (in einem gewissen Sinn, der den wahrscheinlichkeitstheoretischen Begriff verallgemeinert) korreliert sind, Ergebnisse in Form der aus der Physik bekannten Bellschen Ungleichungen finden (siehe [3], [17]) oder Kriterien angeben, die sich aus der Eigenschaft ableiten, dass in der klassischen Mechanik je zwei Observable kommutieren, was bedeutet, dass die zugehörigen Logiken Boolesche Algebren sind.

Im Kontext von numerischen Ereignissen bedeutet die Eigenschaft, dass zwei S-Wahrscheinlichkeiten p_1 und p_2 kommutieren, wofür wir $p_1 C p_2$ schreiben, dass p_1 und p_2 in einer Booleschen Unteralgebra einer (unbekannten) Algebra von S-Wahrscheinlichkeiten liegen, d.h. die Menge $\{p_1, p_2\}$ in eine Boolesche Unteralgebra von W einbettbar ist. Wie man zeigen kann, ist dies auch dadurch charakterisierbar, dass es in W drei weitere numerische Ereignisse x_1, x_2, x_{12} gibt, sodass $p_1 = x_1 + x_{12}$, $p_2 = x_{12} + x_2$ und $x_1 + x_2 + x_{12} \leq 1$ ist. Die beschriebene Situation lässt sich auf eine beliebige Menge $\{p_1, p_2, \ldots, p_n\}$ von S-Wahrscheinlichkeiten einer (unbekannten) Algebra von S-Wahrscheinlichkeiten W verallgemeinern; im Fall $n = 3$ erhält man dabei:

Satz 4.2. *([4]) Eine Menge $\{p_1, p_2, p_3\}$ von S-Wahrscheinlichkeiten von W lässt sich genau dann in eine Boolesche Unteralgebra von W einbetten, wenn $x_1, x_2, x_3, x_{12}, x_{13}, x_{23}, x_{123}$ in W existieren, sodass $x_1 + x_{12} + x_{13} + x_{123} = p_1$, $x_2 + x_{12} + x_{23} + x_{123} = p_2$, $x_3 + x_{13} + x_{23} + x_{123} = p_3$ und $x_1 + x_2 + x_3 + x_{12} + x_{13} + x_{23} + x_{123} \leq 1$.*

Die in Satz 4.2 und dessen Verallgemeinerung auftretende Problematik der Suche nach zahlreichen S-Wahrscheinlichkeiten kann durch eine rekursive Vorgangsweise reduziert werden. Diese Möglichkeit beruht auf der Einführung einer Relation $C(a)$, welche durch „es gilt $p_1 C(a) p_2$ für $p_1, p_2, a \in W$, falls $a \leq p_1 \leq a + p_2 \leq 1$“, definiert ist. Wie man zeigen kann, ist $p_1 C(a) p_2$ äquivalent zu $p_1 C p_2$, was wiederum heißt, dass die Menge $\{p_1, p_2\}$ in eine Boolesche Unteralgebra von W einbettbar ist. Angenommen, für p_1, p_2 gilt $p_1 C(a) p_2$, so ist $a = p_1 \wedge p_2'$ und $p_1 - a = p \wedge p_2$, d.h., a und auch $p_1 - a$ sind mittels der gegebenen S-Wahrscheinlichkeiten darstellbar. Analoges gilt für drei S-Wahrscheinlichkeiten und kann darüber hinaus für n, $n \in \mathbb{N}$ Elemente formuliert werden. Für $n = 3$ erhält man:

Satz 4.3. *([10]) Die Menge $\{p_1, p_2, p_3\} \subseteq W$ ist genau dann in eine Boolesche*

Unteralgebra von W einbettbar, falls $a_{ij} := p_1 \wedge p_j'$ *für* $i < j, i, j \in \{1,2,3\}$ *und* $a_{1213} := p_1 \wedge p_2 \wedge p_3'$ *in W existieren und für diese gilt:* $p_i C(a_{ij})\, p_j$ *für* $i < j, i, j \in \{1,2,3\}$, $(p_1 - a_{12})C(a_{1213})\,(p_1 - a_{13})$ *und* $(p_1 - a_{12})C(a_{1213})\,(p_2 - a_{23})$.

Es ist naheliegend, dass das Konzept der Algebren von *S*-Wahrscheinlichkeiten dadurch verallgemeinert werden kann, dass man weniger einschneidende oder auch etwas andersgeartete Forderungen an Halbordnungen von numerischen Ereignissen stellt (vgl. etwa [8],[11]) und dass es unzählige Resultate gibt, welche in direktem oder indirektem Zusammenhang zu den vorgestellten Ergebnissen über Quantenlogiken stehen. Siehe hierzu etwa die Monografie [18] oder, als kurze Einführung mit einer reichhaltigen Bibliografie, die Übersichtsartikel [19] und [20].

Literatur

[1] J. S. Bell, Speakable and Unspeakable in Quantum Mechanics, 2. Aufl., Cambridge University Press, Cambridge 2004,

[2] E. G. Beltrametti and G. Cassinelli, The Logic of Quantum Mechanics. Addison-Wesley, Reading, Mass., 1981.

[3] E. G. Beltrametti and M. J. Mączyński, On a characterization of classical and nonclassical probabilities. J. Math. Phys. **32** (1991), 1280-1286.

[4] E. G. Beltrametti and M. J. Maczyński, On the characterization of probabilities: A generalization of Bell's inequalities. J. Math. Phys. **34** (1993), 4919-4929.

[5] E. G. Beltrametti, D. Dorninger and M. J. Maczyński, On a cryptographical characterization of classical and nonclassical event systems. Rep. Math. Phys. **60** (2007), 117-123.

[6] G. Birkhoff and J. von Neumann, The Logic of Quantum Mechanics. Annals of Mathematics, Second Series **37**(1936), 823-843

[7] G. Dorfer, D. Dorninger and H. Länger, On the structure of numerical event spaces. Kybernetica **46** (2010), 971-981.

[8] D. Dorninger, On the structure of generalized fields of events. Contr. General Algebra **20** (2012), 29-34.

[9] D. Dorninger and H. Länger, Testing for classicality of a physical system. Intern. J. Theor. Phys. **52** (2013), 1141-1147.

[10] D. Dorninger and H. Länger, Probability measurements characterizing the classicality of a physical system. Rep. Math. Phys. **73** (2014), 127-135.

[11] D. Dorninger and H. Länger, On bounded posets arising from quantum mechanical measurements. Intern. J. Theor. Phys. **55** (2016), 4453-4461.

[12] D. Dorninger and H. Länger, Structural properties of algebras of *S*-probabilities. Math. Slovaca **68** (2018), 485-490.

[13] R. M. Godowski, Commutativity in orthomodular posets. Rep. Math. Phys. **18** (1980), 347-351.

[14] G. W. Mackey, Mathematical Foundations of Quantum Mechanics, Dover Publ., Mineola, NY, 2004.

[15] M. J. Mączyński and T. Traczyk, A characterization of orthomodular partially ordered sets admitting a full set of states. Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys. **21** (1973), 3-8.
[16] P. Pták, Concrete quantum logics. Intern. J. Theor. Phys. **39** (2000), 827-837.
[17] I. Pitowsky, Quantum Probability - Quantum Logic, Lecture Notes in Physics, Vol. 321. Springer, Berlin, 1989.
[18] K. Svozil, Quantum logic, Springer, Series Discrete Mathematics and Theoretical Computer Science, Singapore, 1998.
[19] K. Svozil, Quantum logic. A brief outline. Cornell University, arXiv:quant-ph/9902042v2, 2005.
[20] A. Wilce, Quantum Logic and Probability Theory, The Stanford Encyclopedia of Philosophy, 2017.

Adresse des Autors:
Dietmar Dorninger
Technische Universität Wien
Institut für Diskrete Mathematik und Geometrie
Wiedner Hauptstraße 8-10/104
1040 Wien
email dietmar.dorninger@tuwien.ac.at

Internat. Math. Nachrichten
Nr. 240 (2019), 11–29

Learning how to Listen: Time-Frequency Analysis meets Convolutional Neural Networks

Monika Dörfler
Universität Wien

In this article, the central ideas of convolutional neural networks, a particularly successful class of architectures in deep learning, are introduced. For audio signals, pre-processing data prior to input to network architectures is standard. The interplay between the pre-processing steps, also known as feature extractors, and actual neural network is investigated and the results are illustrated by some applications.

1 Introduction and Motivation

Methods and results from artificial intelligence and data science in general and from deep learning (DL) in particular have had a lot of media presence in the past decade. This attention is due to the sometimes surprising successes of new DL methods in various applications. Some applications of DL in image and video processing include colorisation of black and white images, adding sounds to silent movies, object classification in photographs or automatic game playing (Checkers, Go). Concerning audio signals, mainly for problems in speech, language and music processing, successful applications include speech recognition, singing voice detection, speech and music segmentation or acoustic monitoring of biodiversity[1]. Many other applications e.g. in medicine, finance, mobility (e.g. automated driving) have equally been reported on – DL seems to be able "to do anything".

In this contribution we try to describe the main ideas and mechanisms of a particularly successful kind of network architectures, namely of *convolutional neural*

[1] http://amoby.ofai.at

ISSN 0020-7926

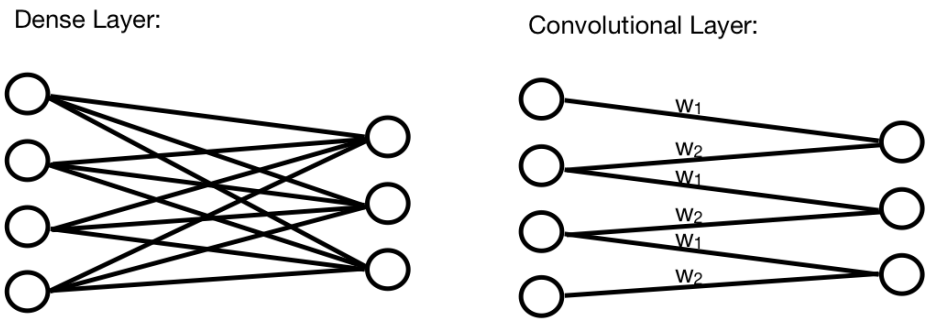

Figure 1: Schematic comparison between a dense and a convolutional layer. In dense layers, each input node is connected with each output node and all weights w_{ij} may be different. In convolutional layers, connectivity is sparse (locality) and weights are shared among nodes.

networks (CNNs). A neural network consists of an input layer, given by the components of the input vector $x \in \mathbb{R}^L$, one or several hidden layers and the output layer. Nodes of subsequent layers are connected by weights, by which the components of an output vector of the previous layer are multiplied, summed and sent through a non-linearity, see (4).

In many cases, the non-linearity step is followed by a dimension-reduction called *pooling*, for which $x_n \in \mathbb{R}^{d(n)}$, the output of layer n, is partitioned in overlapping or non-overlapping sub-vectors or -arrays and mapped to the vector or array of norms of each of the sub-vectors, see [7, Section 3.1] for technical details. The most common approach used in pooling is max pooling, in which case the norm is the maximum norm. The weights connecting the nodes in a given layer (either input or hidden layer) to the nodes of the subsequent layer (either hidden or output layer) define a linear mapping from $\mathbb{R}^{d(n)}$ to $\mathbb{R}^{d(n+1)}$, which, by adding an optional bias b_n, becomes an affine mapping, see (4). The linear mapping in several or all network layers in CNNs is restricted to a special form, namely a block-Toeplitz matrix. In other words, the convolutional layers comprise a certain number of convolutions usually defined by few filter coefficients. The basic structure of a dense layer and a convolutional layer is depicted schematically in Figure 1, see [8, Section 6 and 9] for much more detailed illustrations.

In general, deep learning is based on neural network architectures, with "deep" alluding to the fact that several, in particular more than one and often many, hidden layers are used in the design of the architectures. This depth, rather than "width" in the sense of just using many nodes in a single hidden layer, has been shown to exponentially increase the complexity, or expressivity, of the network, such that increasingly difficult problems can be tackled, see for example [22]. In addition,

the idea of locality and weight sharing in the convolutional layers allows for efficient extraction of certain local features in the data. Once a particular architecture has been chosen, all parameters (or weights) which determine the actual mapping from data to targets, are learned using large sets of labeled data and gradient descent-based weight inference.

Two of the main reasons for the *recent* success of DL are the ability of modern GPUs to process massive amount of data *and* the availability of the necessary size of data bases. Additionally, some new "tricks" have significantly boosted the performance of learning algorithms in past years, for example drop-out [21], stochastic gradient descent [5] or data augmentation [18].

CNNs were originally introduced in image processing problems [12] and subsequently, the core ideas were transferred to audio processing. Since CNNs act on two- or three-dimensional arrays rather than one-dimensional time-series, audio input signals almost always undergo pre-processing in order to map them into an image, which represents a *time-frequency* representation of the respective signal. It turns out, that the time-frequency image of an audio signal encodes signal structures which are essential to human hearing and can thus be expected to be relevant for most learning problems. We will see in detail how this property of the signal representation leads to the possibility to work with significantly less data while achieving the same level of performance.

From a mathematical point of view, whenever one seeks to find representations of a signal (or function), which highlights particular properties, methods from (applied) harmonic analysis come into play. (Applied) harmonic analysis studies representation of functions (signals) as superposition of some basic waves. While in some cases, the collection of these basic building blocks may form a basis, *frames* provide a more flexible concept for signal representation.

Definition 1 (Frames)**.** *Let $\mathfrak{H}$ be a Hilbert space. A sequence $\{g_j : j \in J\} \subseteq \mathfrak{H}$ is called frame, if there exist $A, B > 0$ such that $\forall f \in \mathfrak{H}$*

$$A\|f\|^2 \leq \sum_{j \in J} |\langle f, g_j \rangle|^2 \leq B\|f\|^2. \tag{1}$$

Then $\forall f \in \mathfrak{H}$

$$f = \sum_{j \in J} \langle f, g_j \rangle \tilde{g}_j \tag{2}$$

for a so-called dual frame $\tilde{g}_j$.

We will introduce special frames in Section 3.1 and see how they allow for structured and flexible representations of time-series such as audio signals. The understanding of the interaction of signal representation and network architectures in learning problems is at the core of this article. We start by introducing the main concepts of learning in the next section.

2 Learning as a mathematical paradigm

The process of learning, as discussed in the introduction, may be cast into mathematical formalism as follows. We assume the important case of *learning from data*. In this case, we are given a data set $\mathcal{D} \subset \mathcal{X}$ in an input space $\mathcal{X}$, together with some information about the data, often called "annotation", which, mathematically speaking, is given in the output space and denoted by $\mathcal{T} \subset \mathcal{Y}$. Learning the relationship between $\mathcal{D}$ and their annotations in $\mathcal{Y}$ can then be understood as looking for a function $f : \mathcal{X} \mapsto \mathcal{Y}$, which describes with sufficient accuracy the "nature of data" in relation to the given annotation, [14]. Two simple and important examples, which are easy to interpret mathematically, are regression and classification. Regression is probably the most basic and commonly used type of predictive analysis. It predicts values of a desired target quantity when the target quantity is continuous and depends on the values of one or several independent variables. Therefore, in this setting, we have $\mathcal{X} = \mathbb{R}^d, \mathcal{Y} = \mathbb{R}$.

The second important and widely common application is classification, which pertains to the process of predicting the class, label or target of a given data point. In this case, the dependent variable is categorical, and hence $\mathcal{X} = \mathbb{R}^d, \mathcal{Y} = \{c_1, \ldots, c_n\}$, $c_j \in \mathbb{R}$.

The goal of learning is to find a function f which approximates the mapping between input data and targets sufficiently well, such that for new input data $x \in \mathcal{X}$, correct output variables, or targets, in $\mathcal{Y}$ are predicted by f.

For many learning problems, it is reasonable to start with certain assumptions about the function we aim at finding for the approximation of the given data. In the case of *parametric* modelling, the class of functions is determined by a certain number, often finite, of parameters, whose choice then defines the actual function f. One of the simplest example is *polynomial regression*, which we use here as a showcase for some important notions. In polynomial regression, the relationship between the independent variable x and the dependent variable y is modelled as an n-th degree polynomial in x. Here the basic building blocks are simple monomials, i.e. $\varphi_0 = 1$, $\varphi_1 = x$, $\varphi_2 = x^2$ and so on, up to the chosen maximal degree of the polynomial. The function f is therefore chosen to be a member of the space of polynomials of degree up to n.

The space $\mathcal{H}$, from which the function eventually modeling the data is chosen, is called *hypothesis class*, hence, in the current case, with $\varphi_j = x^j$, we have $\mathcal{H} = P_n$, the vector space of polynomials of degree up to n:

$$\mathcal{H} = \{f : f = \sum_{j=0}^{n} c_j \varphi_j, c_j \in \mathbb{R}\}. \tag{3}$$

For a given set of data, it may not be straightforward to choose the optimal degree of a polynomial which is used to fit the given data and predict the values of

unseen data. The effect of choosing a low degree of freedoms may result in bad approximation even of known data. On the other hand, choosing the degree too high bears the risk of overfitting, that is, small effects that are due to measurement errors or other sources of noise are fitted by the function and then lead to increased error rates on validation on new, unseen data. The quality of approximation will also depend on the number of given data points; naturally, the degree of the polynomial should not be higher than the size of the given data set and will usually be much lower. Finally, one needs to think about how to estimate the coefficients c_j, $j = 0, \ldots, n$. Since polynomial regression under the assumption of zero-mean white noise leads to a linear least-squares problem, it has a closed-form, unique solution, cp. [16]. Estimating the parameters which define a function within more complex hypothesis classes, such as CNNs introduced in the next section, can be much more difficult in general. However, considering the nature of polynomial regression for data estimation, we can identify three central questions:

1. What is the nature of the function f we need to learn?
   Answering this question leads to the choice of a model or hypothesis class $\mathcal{H}$. Problems related to the quality of this choice can be understood as approximation problems: how well can the function underlying the given data in principle be approximated within the chosen hypothesis class.

2. Can I determine a close approximation of the "best" function $f_{\mathcal{H}} \in \mathcal{H}$ using information from the available data points?
   Answering this question leads to the problem of finding optimal parameters within the previously chosen hypothesis class. Questions on how to do so and how to choose optimization algorithms are related to sampling problems.

3. How will the learned function f perform on unseen data?
   The last question is of utmost importance, since it decides about the quality of the learned function with respect to new data, naturally the most interesting aspect, namely its generalisation properties.

In the next section, the hypothesis class of CNNs is introduced.

2.1 Structure of Convolutional Neural Networks

We now introduce the technical details of CNNs, which are in the center of this contribution.

The most basic building block of layer n in a general neural network, compare Figure 1, may be written as

$$x_{n+1} = \sigma(A_n x_n + b_n), \text{ where} \tag{4}$$

- $x_n \in \mathbb{R}^{d(n)}$ is the data vector or array in the n-th layer and b_n is the vector of biases in the n-th layer. A_n is a matrix containing the weights connecting the nodes in the n-th layer to those of the subsequent one, i.e. $A_n(i,j) = w_{ij}$.

- $\sigma : \mathbb{R} \to \mathbb{R}$ is a non-linear function, which is, by a slight abuse of notation, applied component wise. Typcial examples include sigmoidal function, i.e. $\sigma(x) = \frac{1}{1+e^x}$, the rectified linear unit $ReLU(x) = \max(0,x)$, the modulus or the modulus squared.

The basic, modular structure of CNNs has often been described, see e.g. [8]. Note that in each layer the array x_n may have a different dimension. In the case of convolutional layers of CNNs, the matrices A_n have a particular structure for the convolutional layers, namely, it is a block-Toeplitz matrix, or, depending on the implementation of the convolution, a concatenation of circular matrices, each representing one convolution kernel with, referring again to Figure 1, k non-zero coefficients $w_1, \ldots, w_k$. There may be an arbitrarily high number of convolutional layers, followed by a certain number of so-called dense layers, for which A_n is again an arbitrary linear operator. For many applications in audio processing, the chosen architecture comprises four to eight convolutional and two or three dense layers, cp. [9]. Often the first part of the CNN, comprised of the convolutional layers, is denoted as the feature learning stage, while the final part of the network, comprised of at least one, and often several, dense layers, is called classification stage.

2.2 The notion of features

In machine learning the choice of *features* (or data representation), cf. [4], heavily influences the results of learning models and algorithms. The notion itself is not well-defined, mathematically one may assume that good features reduce the dimensionality while maintaining the important part of information contained in the data. The important part of information, in turn, depends on the actual problem to be learned. Therefore, while to some extent feature engineering can be useful, learned feature extractors seem to be able to tune the variability of the extracted, low-dimensional data representation directly to the problem at hand.

Formally, one can say that a feature extractor $\Phi = (\Phi_k)_{k=1}^d : \mathbb{R}^L \mapsto \mathbb{R}^{M_1 \times \ldots \times M_d}$ aims at allowing for a decomposition $f(x) = f_0(\Phi(x))$, where f_0 is (much) simpler than f. The ideally simple f_0 would be given by a linear classifier:

We say that Φ *separates* f *linearly*, if $f(x)$ is well-approximated by a one-dimensional projection:

$$\tilde{f}(x) = \langle \Phi(x), w \rangle = \sum_{k=1}^{d} w_k \cdot \Phi_k(x). \tag{5}$$

As in common linear regression, the vector w in (5) is optimised by minimizing a loss or error term on the training data.

2.3 Loss function and learning of parameters

A loss function is used both as a basis for training a model, that is, for learning the optimal parameters for the description of the desired mapping from input to output data using a function from $\mathcal{H}$ in (3), and for measuring the performance of the model once a parameter vector $\theta = \{c_j, j = 0, \ldots, n\}$ has been learned. One assumes that data $\mathcal{Z}_m = \mathcal{D} \times \mathcal{T} = \{(x_1, y_1), \ldots, (x_m, y_m)\}$ are drawn i.i.d. from a (usually unknown) probability density $\rho : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$. For any function $f : \mathcal{X} \mapsto \mathcal{Y}$, define expected loss (or "true error") as

$$E_\rho(f) := \int_{\mathcal{X}\times\mathcal{Y}} \ell(f, x, y) d\rho(x, y).$$

Further, given a hypothesis space $\mathcal{H}$, parametrised by θ, and a set of annotated data $\mathcal{Z}_m$, the empirical loss function $E_{\mathcal{Z}_m}$ of f_θ with respect to $\mathcal{Z}_m$ is defined as

$$E_{\mathcal{Z}_m}(f_\theta) = \frac{1}{m} \sum_{i=1}^{m} \ell(f_\theta, x_i, y_i).$$

Common, important examples of loss functions include the quadratic loss $\ell(\theta, x_i, y_i) = (f_\theta(x_i) - y_i)^2$, where f_θ is the function determined by learned parameter vector θ, and the categorical cross-entropy loss. The latter is the concatenation of the softmax function on the output vector $\hat{\mathbf{y}} = (f_\theta(x_1), \ldots, f_\theta(x_m))$ and the cross-entropy loss; in other words, in the case of categorical cross-entropy loss, we have

$$\ell(\theta, x_i, y_i) = -y_i \log \frac{e^{f_\theta(x_i)}}{\sum_{j=1}^{m} e^{f_\theta(x_j)}}.$$

If we allow f to be taken from a very large hypothesis class, we will find a function that leads to a rather small error on the given data set, but may generalise badly to unseen data due to overfitting. For example, a lookup-table for the training set would give excellent results for that set, but contain no information about any other points: if we make no restriction on the class of functions from which we choose our estimate, we cannot hope to learn anything. On the other hand, if $\mathcal{H}$ is chosen wisely, in accordance with the assumed or observed properties of $\mathcal{Z}_m$, this a priori choice may have a positive, regularizing effect on the learning process.[2]

In summary, we may try to restrict the complexity of the hypothesis class for f_θ which should ultimately describe the mapping $x_i \to y_i$ by exploiting any prior

[2] Apart from choosing an appropriate hypothesis class, other regularisers play an important role in the optimization process, among others drop-out or weight regularization.

knowledge we may have about the nature of the data or the problem to be learned. Information about the data may be captured by using a fixed feature extractor. In particular, the feature extractor should be designed in such a way as to encode known symmetries in the data which imply invariance of the mapping with respect to these symmetries. The interplay of data properties, features and network architectures is studied in the next section with a focus on representation of audio data.

3 Features and Network Design

In order to give an idea of the possible interactions between (the amount of available) data, feature extraction and network design, we briefly discuss two examples.

In the recent contribution [15], a deep CNN learns semantic music content from raw audio data with up to more than 90% accuracy. The authors compare the performance of the network that acts directly on the audio waveform by applying 7 subsequent one-dimensional convolutional layers, consisting of 64 to 256 filters each, to the performance of a network with a single two-dimensional convolutional layer acting on the audio signal's mel-spectrogram by applying 480 convolutional filters. The mel-spectrogram, as we shall see in Section 3.1, provides a rather structured representation of audio signals, which also incorporates, by averaging over frequency, invariances known to be important to the human hearing system. The important observation in [15] is that the network based on mel-spectrogram input always yields better results than the network based on raw waveforms, except if tremendously large amounts of data, more precisely, more than half a million songs, or around 40.000 hours of music, are available for learning.

On the other hand, in [7], it was shown that allowing for a controlled adaptivity in the design of time- and frequency-averaging can slightly, but significantly, improve learning results, cf. Section 3.4.2 for details.

We next review some classical time-frequency representations.

3.1 Spectrogram, Mel-Spectrogram and Gabor Frames

The standard input in learning methods for audio signal is based on a sub-sampled spectrogram, either in its raw form, or after some pre-processing in the form of averaging over time or frequency, such as the computation of mel-spectrogram, defined in (7). Mathematically, the samples of a spectrogram may be understood as the coefficients of a function with respect to the members of a particular frame, recall Definition 1. Gabor frames are given as a collection of time- and frequency

shifted versions of a basic building block or window g, where $T_\xi g(t) = g(t-\xi)$ denotes a time-shift by ξ and $M_\eta g(t) = g(t) \cdot e^{2\pi i \eta t}$ denotes a frequency-shift, or modulation, by η.

Example 1 (Gabor frame)**.** *A set of functions* $\mathcal{G}(g,\alpha,\beta) = \{g_{k,l} = M_{\beta k} T_{\alpha l} g : k,l \in \mathbb{Z}\}$ *is called Gabor frame, if it satisfies the frame inequality* (1) *for some constants* $0 < A, B < \infty$.

In order to distinguish between the function f, which denotes the mapping $\mathcal{D} \to \mathcal{T}$, and the data points, which may also be described as functions, signals or vectors, we now think of data points as $x \in \mathfrak{H}$, where $\mathfrak{H}$ can be any Hilbert space of functions, sequences or vectors.

Note that the short-time Fourier transform (STFT) of x with respect to a time-localised window g, e.g. a Gaussian or a Hanning window, gives the time-localised frequency content of x. This function is at the heart of time-frequency analysis, cf. [10], and may be defined as follows, with $\mathcal{F}$ denoting Fourier transform:

$$\mathcal{V}_g x(\xi,\eta) = \mathcal{F}(x \cdot T_\xi g)(k) = \int_t x(t) g(t-\xi) e^{-2\pi i \eta t} dt. \tag{6}$$

The spectrogram S_x of x then consists of the absolute value squared of samples of $\mathcal{V}_g x$ on a rectangular lattice $\alpha\mathbb{Z} \times \beta\mathbb{Z}$, i.e.: $S_x(l,k) = |\mathcal{V}_g x(\alpha l, \beta k)|^2 = |\langle x, g_{k,l}\rangle|^2$. From the comparison between raw audio waveform $x(t)$ and its spectrogram S_x in Figure 2 it becomes obvious, that S_x expresses essential signal properties much more clearly, or sparsely, than raw audio data. In particular, the variations of the energy over time and frequency are separated, which corresponds to the human perception of sound.

Furthermore, time-frequency representations intrinsically introduce approximate invariance to small, local time-frequency modifications and, in the case of spectrogram, to phase-shifts. Further invariance can by introduced by averaging over computed coefficients. The widely used mel-spectrogram introduces invariance to small frequency-shifts by averaging S_x according to a filterbank, which approximately mimics the behaviour of the human ear:

Example 2 (Mel-spectrogram, [20])**.** *The mel-spectrogram is derived from* S_x *by taking weighted averages over frequency channels defined by the mel-scale:*

$$\mathrm{MS}_g(x)(l,\nu) = \sum_k S_x(l,k) \cdot \Lambda_\nu(k). \tag{7}$$

Figure 3 shows a comparison of spectrogram and mel-spectrogram; the spectrogram has 2048×176, while the mel-spectrogram has only 100×176 entries, due to the averaging over frequency channels. The frequency-information is thus condensed in fewer coefficients and averaging happens in a non-linear manner. In

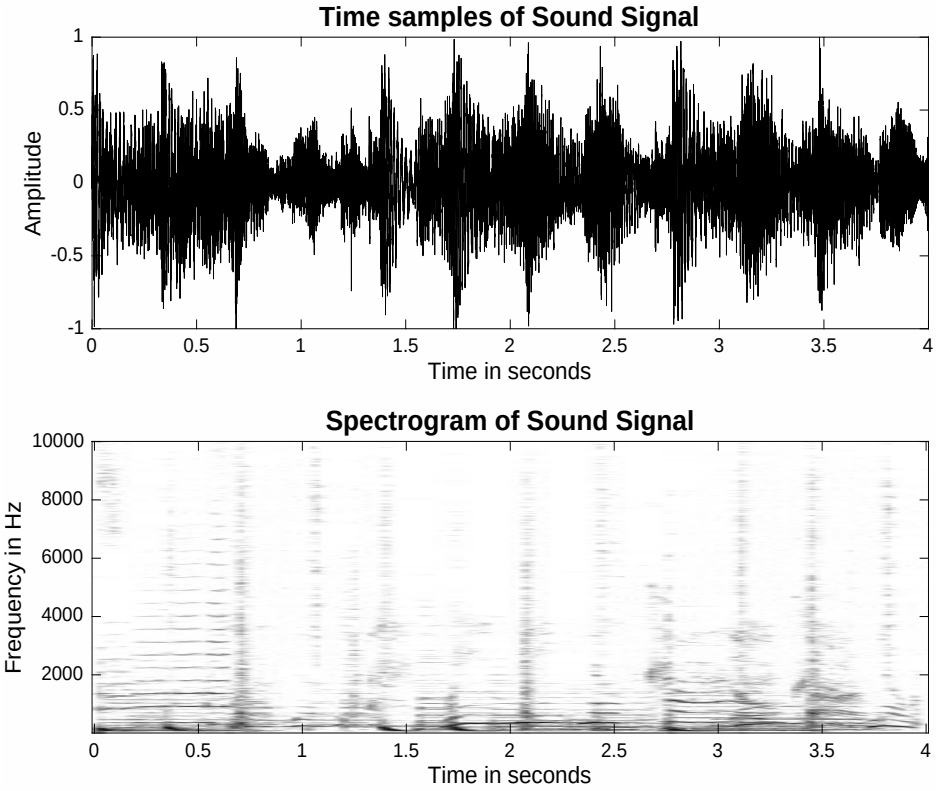

Figure 2: Sound Signal: Raw Time Series (Upper Plot), Spectrogram (Lower Plot); in the visualisation of the spectrogram, higher grey level corresponds to higher amplitude.

fact, linear sampling in frequency leads to an accumulation of most energy contained in an audio signal in lower frequency channels, as can be seen in the lower plot of Figure 3, and this uneven distribution of frequency is remedied by the mel-averaging. Using adaptive frames in the primary analysis of an audio signal x, leads to a more flexible situation for introducing local invariance. The corresponding tool is given by *non-stationary Gabor frames*, [11], which use windows with adaptive bandwidth instead of modulated versions of a fixed window g. We denote by $\mathbf{h}$ the vector of window functions h_ν, where $\nu \in I$ are the frequency channels, and by β the vector collecting the corresponding frequency-shift parameters.

Example 3 (Non-stationary Gabor frame)**.** *A frame given by* $\mathcal{G}_n(\mathbf{h},\beta) = \{h_{\nu,l} = T_{l\beta_\nu}h_\nu : l \in \mathbb{Z}, \nu \in I\}$ *is called non-stationary Gabor frame.*

The adaptive spectrogram S_a contains the coefficients of x with respect to the non-stationary Gabor frame $\mathcal{G}_n$: $S_a(l,k) = |\langle x, T_{l\beta_\nu}h_\nu\rangle|^2$.

Relaxing certain parameters in the feature extractor, as to be able to achieve more flexibility in the representation of the input signals, leads to a situation which

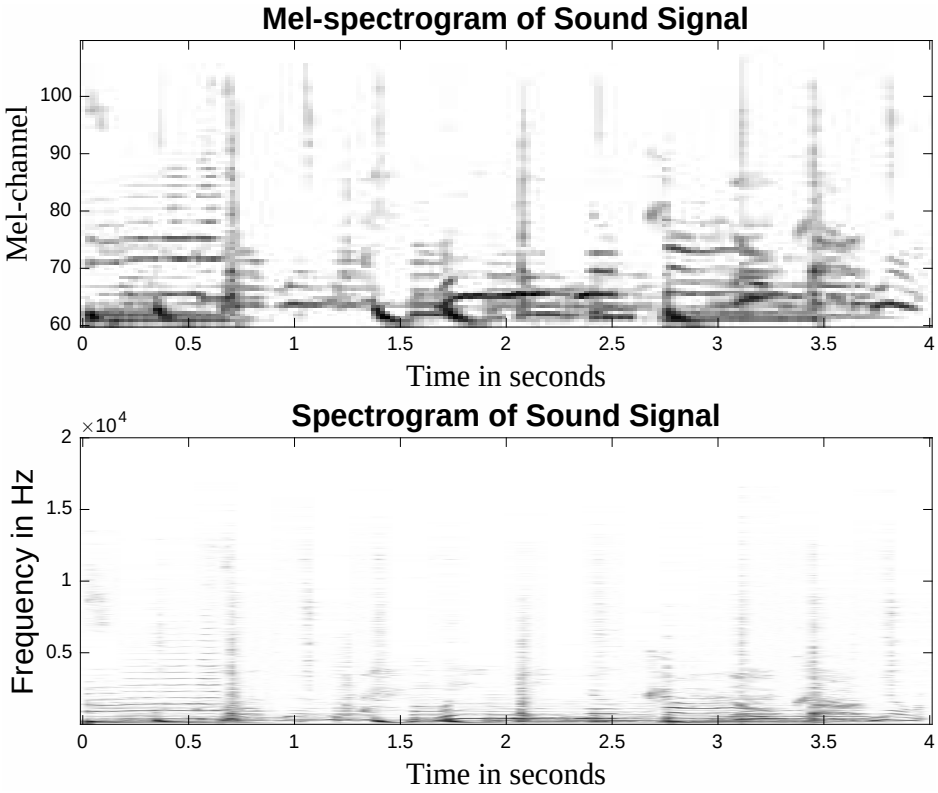

Figure 3: Sound Signal: Mel-Spectrogram (Upper Plot), Spectrogram (Lower Plot)

potentially introduces more variability in the data. Thus, this addition of degrees of freedom needs to be balanced by tools which reduce the variance in a manner that is beneficial for the problem to be solved. We give a formal approach to invariance and stability in the next section.

3.2 Invariance and stability

In some sense, identifying invariance in the representatives of a certain class in a data set is at the core of learning. This is most easily seen when thinking about object recognition in images: obviously the objects, corresponding e.g. to the class "flower", will show a high variability in many aspects and the learner needs to be able to ignore most of the variability in order to learn a function that does not overfit the problem. Similarly, in audio problems, if we want the network to recognise whether or not a human voice is present in a certain part of an audio signal, cp. Section 3.4.2, the network has to learn to ignore a lot of information which is contained in the signal. Some invariances are trivial and global, for example, if an image is rotated, it is still the same image and the network should be globally invariant to the action of the rotation group. Other invariances are more

subtle, and only local, for example small frequency-shifts of some components in a sound may not change its over-all perception. We are led to the following definition.

Definition 2. *Given an augmentation $\mathcal{A}$, that is a set of bounded operators acting on $\mathcal{X}$: $\mathcal{A} = \{T_p : \mathcal{X} \to \mathcal{X}\}$, then f is said to be invariant to $\mathcal{A}$ with respect to $\mathcal{D} \subset \mathcal{X}$, if $f(T_p(x)) = f(x)$ for all $x \in \mathcal{D}$. If $\mathcal{A}$ is parametrised by a set $\mathcal{P}$, on which a metric $|\cdot|_{\mathcal{P}}$ is defined, then we say that f is locally stable to $\mathcal{A}$, if $\|f(T_p(x)) - f(x)\| \leq C \cdot |p|_{\mathcal{P}} \cdot \|x\|$ for all $x \in \mathcal{D}$, all $p \in \mathcal{P}$ and some constant C.*

Note that for categorical problems local stability actually implies local invariance.

Example 4. *Let $\mathcal{D}$ be a data set of audio signals and let $\mathcal{A}(\mathcal{D})$ denote the augmented data-set achieved by multiplication with a phase factor, i.e. $\mathcal{A} = \{T_p : \mathcal{X} \to \mathcal{X}, T_p(x)(t) = x(t) \cdot e^{\pi i p}, p \in \mathbb{R}\}$. Then, a learning problem related to human hearing can be expected to be invariant to phase shifts[3] $\mathcal{A}$ and hence in this case we may factorise $f(x) = f_0(S_x)$.*

The above example explains why using spectrograms of sounds as input to CNNs instead of raw sounds is beneficial for most problems. This may be explained by the fact that explicitly encoding known invariances reduces the generalisation error, as we shall see in more detail in Section 3.3. Computing the spectrogram of an input sound x means to separate time- and frequency-variability in a first step, and to reduce variability due to phase changes in a second step, by taking absolute value of the STFT-coefficients. This basic idea of applying some kind of time-variant filter bank followed by a reduction of variability is at the core of CNNs. For special function classes, such as audio signals, one can introduce more structured layers, as was suggested by Mallat in the form of a scattering transform, cf. [1]. A similar idea based on Gabor frame leads to approximate invariance to certain modifications of audio signals.

Example 5 (Gabor Scattering (GS))**.** *Gabor Scattering (GS), introduced in [3], iteratively applies Gabor transforms with different subsampling schemes, a non-linearity and subsequent time-averaging. The ℓ-th layer of a GS transform is thus defined by*

$$x_\ell = \sigma_\ell(\langle x_{\ell-1}, M_{\beta_\ell j} T_{\alpha_\ell k} g_\ell \rangle_{\mathfrak{H}_{\ell-1}}),$$

where $x_{\ell-1}$ is the output-vector of the previous layer. Time-averaging with ϕ_ℓ yields the corresponding feature extractor:

$$\Phi(f) := \bigcup_{\ell=0}^{\infty} \{x_\ell * \phi_\ell\}. \tag{8}$$

As proved in detail in [3], the output of the first layer in GS captures all frequency

[3]The human hearing system is in general insensitive to phase information, [23].

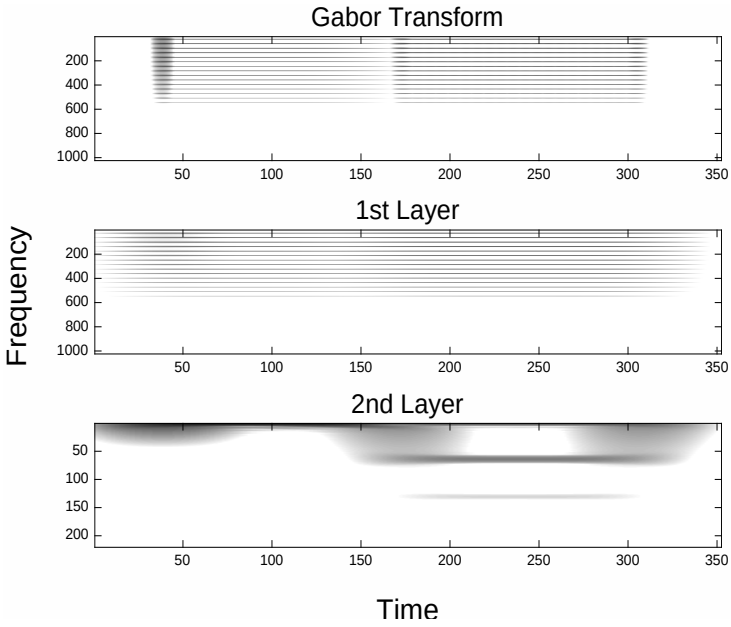

Figure 4: Gabor transform, first and second layer output of Gabor Scattering of a synthetic sound signal. Clearly visible: different invariances induced by first and second layer.

information and is locally stable to amplitude modulations of sound signals. Conversely, the second layer is locally stable to frequency variations and encodes amplitude changes developing over a larger scale of time. This property of GS is illustrated in Figure 4 by means of a synthetic signal comprised of two tones with the same frequency content, but different amplitude-envelope: while the first tone has a constant amplitude, the second one shows amplitude modulation at a certain frequency. This amplitude modulation is clearly present in the second layer, shown in the lowest plot, while the first layer is invariant to this signal property.

Hence, in addition to the separation of variations in time and frequency, provided by the spectrogram, Gabor scattering separates variability with respect to higher-level signal properties. As argued heuristically in [13], separation of meaningful signal-properties will lead to faster and more stable learning of network parameters for certain problems. This hypothesis has been substantiated in [3, 2] on several classification problems for audio signals, cf. Section 3.4.1.

When using feature extractors such as GS, one faces more complex computations prior to the actual learning process. Ideally, this additional computational cost should then allow for more modest architectures used for a particular problem and we hence face the question: can a representation which encodes known symmetries reduce necessary network size?

3.3 Feature-network pairs

Designed feature extractors may encode invariances present in a data set, on the other hand, learning the extent to which a representation should be locally invariant along certain dimensions, may also be useful. Generally, one needs to consider the interaction between feature extractor and network architecture, with respect to a particular set of data.

A CNN $\mathcal{N}$ is understood as the architecture defined by the input dimension, the number of convolutional layers D_c, the number of dense layers D_d, the number and size of convolutional kernels in each convolutional layer and the type of non-linearity and pooling in each layer. We call the parameter vector comprising all the weights occurring in the network by $\theta \in \mathbb{R}^p$ and thus a concrete realisation given a parameter vector θ is denoted by $\mathcal{N}(\theta)$.

The following definition was introduced in [6].

Definition 3 (CNN equivalence)**.** *Given two feature-network pairs $(\Phi_j, \mathcal{N}_j)$, $j = 1,2$, we say that $(\Phi_1, \mathcal{N}_1)$ is subordinate to $(\Phi_2, \mathcal{N}_2)$ with respect to a data set $\mathcal{D}$ and annotations $\mathcal{T}$, if for all $\theta_1 \in \mathbb{R}^{p_1}$ there exists a $\theta_2 \in \mathbb{R}^{p_2}$ such that*

$$\mathcal{N}_1(\theta_1)(\Phi_1(x)) = y \Rightarrow \mathcal{N}_2(\theta_2)(\Phi_2(x)) = y \ \forall (x,y) \in \mathcal{D} \times \mathcal{T}.$$

$(\Phi_1, \mathcal{N}_1)$ and $(\Phi_2, \mathcal{N}_2)$ are equivalent with respect to $\mathcal{D}$, if they are subordinate to each other.

The augmentation of a data set then yields the following result for feature extractors which are invariant under this augmentation.

Proposition 1. *Let $(\Phi_1, \mathcal{N}_1)$ be subordinate to $(\Phi_2, \mathcal{N}_2)$ with respect to $\mathcal{D}$ and let $\mathcal{A}(\mathcal{D})$ denote an augmented data-set.*
If $\mathcal{N}_1(\Phi_1(\mathcal{A}(x))) = \mathcal{N}_1(\Phi_1(x))$ for all $x \in \mathcal{D}$, and Φ_2 is invariant to $\mathcal{A}$, then $(\Phi_1, \mathcal{N}_1)$ is also subordinate to $(\Phi_2, \mathcal{N}_2)$ with respect to $\mathcal{A}(\mathcal{D})$.

The following application of Proposition 1 uses the augmentation by phase multiplication introduced in Example 4.

Example 6. *Let $(Id, \mathcal{N}_1)$ be subordinate to $(S_x, \mathcal{N}_2)$ with respect to $\mathcal{D}$; let $\mathcal{A}(\mathcal{D})$ denote the augmented data-set achieved by multiplication with a phase factor. If $\mathcal{N}_1$ is invariant to $\mathcal{A}$, then, due to invariance of S_x to $\mathcal{A}$, $(Id, \mathcal{N}_1)$ is also subordinate to $(S_x, \mathcal{N}_2)$ with respect to $\mathcal{A}(\mathcal{D})$.*

Definition 4 (Generalisation error)**.** *Let a probability density $\rho : \mathcal{X} \times \mathcal{Y} \mapsto \mathbb{R}$, a hypothesis space $\mathcal{H}$, a set of annotated data $\mathcal{Z}_m = \mathcal{D} \times \mathcal{T} = \{(x_1, y_1), \ldots, (x_m, y_m)\}$, and f_θ, learned from $\mathcal{Z}_m$, be given. Then the generalisation error $L_{\mathcal{Z}_m}(f_\theta)$ is the gap between empirical and expected risk:*

$$L_{\mathcal{Z}_m}(f_\theta) = E_\rho(f_\theta) - E_{\mathcal{Z}_m}(f_\theta).$$

The following result is a direct consequence of Theorem 2 in [17]. We denote by $\mathcal{C}(\mathcal{X},\varepsilon)$ the covering number[4] of a metric space $\mathcal{X}$.

Proposition 2. *Introducing invariance to augmentation $\mathcal{A}$ in a stable learning algorithm leads to a reduction of the generalisation error by a factor proportional to $\mathcal{C}(\mathcal{D})/\mathcal{C}(\mathcal{A}(\mathcal{D}))$.*

Hence, if a feature extractor Φ is invariant to an augmentation $\mathcal{A}$, learning a $\mathcal{A}$-invariant mapping $f : \mathcal{X} \to \mathcal{Y}$ from $\Phi(\mathcal{X})$ reduces the generalisation gap in comparison to learning from $\mathcal{X}$ or $\tilde{\Phi}(\mathcal{X})$, where $\tilde{\Phi}$ is not invariant to $\mathcal{A}$.

In some obvious cases, as seen in the examples given in Section 3.2, desired invariances may be induced in a feature extractor by a proper choice of filters, non-linearites and sub-sampling parameters. On the other hand, learning the filter coefficients of the convolutional layers in CNNs often seems to provide the invariance, which is most beneficial for the data and problem at hand. The disadvantage is the fact that in order to obtain reliable functions f_θ, a sufficient amount of data needs to be available and the network needs to provide sufficient degrees of freedom. Hence, in particular if a restricted number of annotated data points is available, using sophisticated feature extractors can be useful if not necessary.

In the case of time-frequency representations of audio signal, the amount of invariance in time and frequency can be tuned by learning just a few parameters, as we shall see in the next section.

3.3.1 Invariance in Features: Training or Design?

In many audio-processing tasks, including deep learning applied to audio, the mel-spectrogram is used as the standard signal representation. In [7], we showed that mel-spectrogram coefficients can be obtained by using adaptive filters and subsequent time-averaging. We define the time-averaging operator with respect to a family of averaging functions ϖ_ν, $\nu \in I$ and a non-stationary Gabor frame as

$$\mathrm{TA}^h_\varpi(x)(l,\nu) := \sum_{l'} S_a(l',\nu) \cdot \varpi_\nu(l'-l). \tag{9}$$

Proposition 3. *Let an analysis window g and mel-filters Λ_ν be given, for $\nu \in I$. If for each ν, the windows h_ν and time-averaging functions ϖ_ν are chosen such that*

$$V_{h_\nu} h_\nu(\xi,\eta) \cdot \mathcal{F}(\varpi_\nu)(\eta) = V_g g(\xi,\eta) \cdot \mathcal{F}^{-1}(\Lambda_\nu)(\xi), \tag{10}$$

then the mel-spectrogram coefficients $\mathrm{MS}_g(x)(l,\nu)$ can be obtained by time-averaging the filtered signal's absolute value squared, i.e. $\mathrm{MS}_g(x) = \mathrm{TA}^h_\varpi(x)$.

[4] For a metric space $\mathcal{X}$ and some $\varepsilon > 0$, the *covering number* $\mathcal{C}(\mathcal{X},\varepsilon)$ is given by the minimal number $l \in \mathbb{N}$, such that there exist l disks of radius ε, which cover $\mathcal{X}$.

The following statement is an immediate consequence.

Theorem 1. Consider $\mathcal{N}_1$ as a convolutional network with D_c convolutional layers and a network $\mathcal{N}_2$ which is the same as $\mathcal{N}_1$, except that it has an additional convolutional layer, consisting of a finite number of convolutional kernels with sufficient length in time-direction and length 1 in frequency direction, preceding the D_c convolutional layers in $\mathcal{N}_1$. If the windows g, h_ν and the mel-filters Λ_ν are chosen such that (10) holds, then $(\mathrm{MS}_g, \mathcal{N}_1)$ is subordinate to $(S_a, \mathcal{N}_2)$.

The theorem says that using an adaptive filter bank and subsequent time-averaging, with a flexible length of averaging for each frequency channel, leads to a more powerful representation than using mel-spectrogram coefficients as the input to a standard CNN. At the same time, a similar invariance structure in time-frequency is induced by the implicit set-up of the feature extractor: only the amount of averaging in time and frequency, respectively, is learned. This approach was successfully applied to a standard problem in audio processing, namely singing voice detection, see Section 3.4.2, for which it was possible to show that controlled adaptivity indeed improves the performance of the network.

3.4 Application Examples

For illustration of our theoretical results, we use two recent applications.

3.4.1 Sound classification

A classical problem for CNNs in audio is the classification of sounds. We used GS transforms in comparison with spectrogram coefficients in order to understand the invariance encoded by GS. In [3], GS was used as input to a simple CNN in order to classify synthetic data into four classes, namely "no modulation","amplitude modulation", "frequency modulation" and "frequency + amplitude modulation". In [2], GS was applied to the problem of instrument classification on the Good-sounds data set[5]. In both applications, GS led to better performance when compared to both spectrogram or mel-spectrogram. In particular, for functions f_θ learned on smaller training sets, as expected, GS was able to achieve smaller generalisation error, measured empirically as the difference between training error and validation error.

3.4.2 Singing Voice Detection

Singing voice detection is a binary problem about whether a short audio segment contains human voice, see e.g. [19]. The state-of-the-art architecture in [19] uses

[5] https://www.upf.edu/web/mtg/good-sounds

mel-spectrograms of the audio segments as inputs to a CNN with a total number of 1.41 million weights, thereof 91% for the dense layers, and leads to an error rate of less than 7% on unseen data. In [7], we were able to show that allowing for adaptivity in the choice of parameters of the feature extractor applied to the audio segments leads to an over-all performance comparable to the state-of-the-art while drastically reducing the network size. As argued in Section 3.3.1, this size reduction is possible since useful invariances are pre-designed by the feature extractor, while the precise extent of the invariances is tuned by few trainable parameters.

The CNN architectures used in [7] are as follows. In each of the two investigated architectures, four convolutional layers were used, namely: two 3×3 convolutions (32 and 16 kernels) with subsequent *ReLU*-non-linearity, followed by 3×3 non-overlapping max-pooling and two more 3×3 convolutions (32 and 16 kernels) with subsequent *ReLU*-non-linearity, followed by 3×3 non-overlapping max-pooling.

For the classification stage in the dense layers, two different sizes were compared: Network A uses two dense layers of 64 and 16 units, respectively, leading to a total number of 94337 weights, thereof 85% classification stage. In the even smaller Network B the classification stage consists of only one dense layer with 32 units. Network B thus has a total number of 53857 trainable parameters, thereof 73% classification stage. The final dense layer is a single sigmoidal output unit in both cases.

For evaluating the performance of the trained models, we used the 'area over the ROC curve' measure (AOC), where lower AOC indicates a better result.

For mel-spectrogram input, the smaller networks score 6.66% (Network A) and 7.05% (Network B), respectively, compared to 6.74% of the original architecture.

The different adaptive models exhibit very similar performance measures; as an example, for Network B, we achieved AOC $< 6.75\%$, which is significantly better than the spectrogram-based case at AOC $= 7.05\%$. As expected, the effects of adaptivity on the evaluation results are more pronounced for the smaller architecture (Network B), due the higher complexity of the larger classification stage in Network A. In summary, the experiments strengthen the hypothesis that sophisticated feature extractors which encode prior knowledge and some controlled flexibility, are particularly important, if smaller data bases allowing for less complex network architectures are available.

4 Summary and Perspectives

Deep Learning has reached many areas of relevance, both in research and everyday life. In this article, we described the basic concepts and ideas of convolutional

neural networks and tried to clarify the connection between feature design and network architectures, in particular for learning problems on audio signals. We pointed out that for complex problems, satisfactory results require huge amounts of data. Designing smart feature extractors can lead to better performance and in particular to smaller generalisation error with less data. Encoding known invariances seems to play an important role in reducing generalisation error and thus improving performance on unseen (test) data. The exact formal description of the underlying mechanisms, in particular for various kinds of interacting local and approximate invariances, will be the content of future work.

5 Acknowledgements

The research presented in this article has been supported by the Vienna Science and Technology Fund (WWTF) project SALSA (MA14-018). The author wishes to thank Roswitha Bammer, Pavol Harar, Arthur Flexer and Thomas Grill for the fruitful collaboration and for helping to produce insightful illustrations.

References

[1] J. Andén and S. Mallat, Deep scattering spectrum, *IEEE Transactions on Signal Processing* vol. 62, no. 16, pp. 4114–4128, 2014

[2] R. Bammer, A. Breger, P. Harar and M. Dörfler, Machines listening to music: the role of signal representations in learning from music, *preprint, submitted*, available at https://www.univie.ac.at/nuhag-php/home/publication_doerfler.php, 2019

[3] R. Bammer, P. Harar, M. Dörfler, Gabor frames and deep scattering networks in audio processing , *preprint*, arxiv.org/abs/1706.08818, 2019

[4] Y. Bengio, A.C. Courville and P. Vincent, Representation Learning: A Review and New Perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, pp. 1798–1828, 2013

[5] L. Bottou, Large-scale machine learning with stochastic gradient descent, I, *Proceedings of the 19th International Conference on Computational Statistics* , pp. 177–187, Paris, France, 2010

[6] M. Dörfler, R. Bammer and T. Grill, Inside the Spectrogram: Convolutional Neural Networks in Audio Processing, *Proceedings of Sampta 2017*, doi=10.1109/SAMPTA.2017.8024472, 2017

[7] M. Dörfler, T. Grill, R. Bammer and A. Flexer, Basic Filters for Convolutional Neural Networks Applied to Music: Training or Design? *Neural Computing and Applications*, 2018

[8] I. Goodfellow, Y. Bengio and A. Courville, Deep Learning, *MIT Press*, http://www.deeplearningbook.org, 2016

[9] T. Grill and J. Schlüter, J., Music Boundary Detection Using Neural Networks on Combined Features and Two-Level Annotation, *Proceedings of the 16th Interna-*

tional Society for Music Information Retrieval Conference (ISMIR 2015), Malaga, Spain, 2015
[10] K. Gröchenig, *Foundations of Time-Frequency Analysis*. Appl. Numer. Harmon. Anal. Birkhäuser, Boston, MA, 2001
[11] N. Holighaus, M. Dörfler, G. A. Velasco, and T. Grill, A framework for invertible, real-time constant-Q transforms, *IEEE Trans. Audio Speech Lang. Process.*, vol. 21, no. 4, pp. 775–785, 2013
[12] A. Krizhevsky, I. Sutskever and G. Hinton, ImageNet Classification with Deep Convolutional Neural Networks, *Advances in Neural Information Processing Systems 25*, pp. 1097–1105, 2012
[13] S. Mallat, Understanding deep convolutional networks, *Philos Trans A Math Phys Eng Sci.*, 374(2065), 2016
[14] T.M. Mitchell, Machine Learning, *McGraw-Hill Science/Engineering/Math* ISBN: 0070428077, 1997
[15] J. Pons, O. Nieto, M. Prockup, E.M. Schmidt, A.F. Ehmann and X. Serra, End-to-end learning for music audio tagging at scale, *Proceedings of the 19th International Society for Music Information Retrieval Conference, ISMIR*, Paris, France, 2018
[16] P. Royston, Polynomial Regression, In: Armitage P, Colton T eds. *Encyclopedia of Biostatistics*, Wiley, 2005
[17] J. Sokolic, R. Giryes, G. Sapiro and M.Rodrigues, Generalization Error of Invariant Classifiers, *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, pp. 1094–1103, 2017
[18] J. Salamon and J. Bello, Deep Convolutional Neural Networks and Data Augmentation for Environmental Sound Classification, *IEEE Signal Processing Letters*, vol. 24, pp. 279–283, 2017
[19] J. Schlüter and T. Grill, Exploring Data Augmentation for Improved Singing Voice Detection with Neural Networks, *Proceedings of the 16th International Society for Music Information Retrieval Conference*, 2015
[20] S. S. Stevens, A scale for the measurement of the psychological magnitude pitch, *Acoustical Society of America Journal*, vol. 8, 1937
[21] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *The Journal of Machine Learning Research*, 15(1), pp. 1929–1958, 2014
[22] M. Telgarsky, Benefits of depth in neural networks, *29th Annual Conference on Learning Theory,* pp. 1517–1539, 2016
[23] F.L. Wightman, Pitch and Stimulus Fine Structure, *J.Acoust.Soc.Am.*, 54, pp. 397–407, 1973

Author's address:
Faculty of Mathematics,
Oskar-Morgenstern-Platz 1,
A-1090 Wien, AUSTRIA
email monika.doerfler@univie.ac.at

Internat. Math. Nachrichten
Nr. 240 (2019), 31–41

Free Semialgebraic Geometry

Tim Netzer
Universität Innsbruck

This is a survey article on the currently very active research area of free (=non-commutative) real algebra and geometry. We first review some of the important results from the commutative theory, and then explain similarities and differences as well as some important results in the free setup.

1 Classical Semialgebraic Geometry

In this section we briefly review some concepts and results from commutative real algebra and geometry. For details and proofs see for example [23, 18, 3, 19, 20].

Important objects of study in classical (=commutative) real algebra and geometry are semialgebraic sets. A *basic closed semialgebraic set* is of the form

$$W(p_1,\dots,p_r) := \left\{a \in \mathbb{R}^d \mid p_1(a) \geq 0,\dots,p_r(a) \geq 0\right\}$$

where $p_1,\dots,p_r \in \mathbb{R}[x_1,\dots,x_d]$ are polynomials. A general *semialgebraic set* is a (finite) Boolean combination of basic closed semialgebraic sets. An important

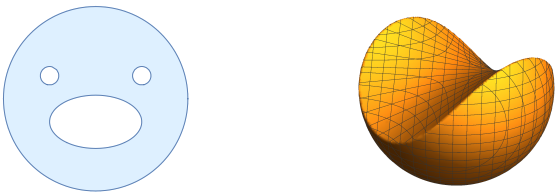

Figure 1: Two (basic closed) semialgebraic sets

and foundational result in real algebraic geometry is the following:

ISSN 0020-7926

Theorem 1 (Projection Theorem)**.** *Projections of semialgebraic sets are again semialgebraic.*

The Projection Theorem is not easy to proof. It also has some strong implications for logic and model theory of real closed fields. Since projections correspond to existential quantifiers, it leads to quantifier elimination in the theory of real closed fields, which lies at the core of almost any Positivstellensatz in real algebra. It also proves decidability of the theory of real closed fields. Although not easy, the Projection Theorem admits constructive proofs. In practice, finding a semialgebraic description of a projection can be very challenging, however.

In classical algebraic geometry, affine varieties are classified via polynomial functions that vanish on them. These functions are described algebraically by Hilbert's Nullstellensatz. For semialgebraic sets, one considers *nonnegative* polynomials, and *Positivstellensätze* provide algebraic characterization of such polynomials. The role of ideals is taken by *preorderings*. The preordering generated by polynomials $p_1,\dots,p_r \in \mathbb{R}[x_1,\dots,x_d]$ arises from the p_i and sums of squares of polynomials, by addition and multiplication:

$$\mathcal{P}(p_1,\dots,p_r) := \left\{ \sum_{e\in\{0,1\}^r} \sigma_e p_1^{e_1}\cdots p_r^{e_r} \mid \sigma_e \text{ sums of squares of polynomials} \right\}.$$

Note that polynomials from the preordering are obviously nonnegative on $W(p_1,\dots,p_r)$.

Theorem 2 (Nichtnegativstellensatz)**.** *For $p, p_1,\dots,p_r \in \mathbb{R}[x_1,\dots,x_d]$, the following are equivalent:*

(i) $p \geq 0$ *on* $W(p_1,\dots,p_r)$.

(ii) $t_1 p = p^{2n} + t_2$ *for some* $t_1, t_2 \in \mathcal{P}(p_1,\dots,p_r)$, $n \in \mathbb{N}$.

The direction *(ii)*⇒*(i)* is straightforward to see, so *(ii)* is an algebraic certificate for nonnegativity of p. The factor t_1 in *(ii)* is often called a *denominator*. Over the field $\mathbb{R}(x_1,\dots,x_d)$ it can be brought to the other side, providing a representation with rational functions. The case $r = 1$ and $p_1 = 1$ is precisely Hilbert's 17th Problem: Every globally nonnegative polynomial is a sum of squares of rational functions. One can get rid of the denominator only under additional assumptions. The first and most important such Positivstellensatz without denominators is the following, where the conditions of boundedness and strict positivity is necessary for the theorem to hold:

Theorem 3 (Schmüdgen's Positivstellensatz)**.** *Let $p_1,\dots,p_r \in \mathbb{R}[x_1,\dots,x_d]$ and assume $W(p_1,\dots,p_r)$ is bounded. Then for any $p \in \mathbb{R}[x_1,\dots,x_d]$*

$$p > 0 \text{ on } W(p_1,\dots,p_r) \quad \Rightarrow \quad p \in \mathcal{P}(p_1,\dots,p_r).$$

After this very brief introduction, let us now pass to the non-commutative setup.

2 Free Real Algebra and Geometry

Semialgebraic sets are defined by polynomial inequalities. So before we can talk about non-commutative semialgebraic sets, we introduce non-commutative polynomials. In the non-commutative setup we will always use an involution (in fact the involution is also there in the classical case, however invisible since it is just the identity on real polynomials). In presence of an involution we can use complex numbers as our ground field and take Hermitian elements as the "real" objects. Using complex numbers is often more convenient and allows for cleaner proofs.

So let $\mathbb{C}\langle z_1,\dots,z_d\rangle$ denote the algebra of polynomials in the *non-commuting variables* $z_1,\dots,z_d$. Their elements are $\mathbb{C}$-linear combinations of *words* in the variables. Since the variables do not commute, words like z_1z_2 and z_2z_1 are different, where they would coincide in the commutative case. We use the involution $*$ on $\mathbb{C}\langle z_1,\dots,z_d\rangle$ that fixes the variables, i.e. $z_i^* = z_i$ holds for all i, but reverses the order in each word and acts as complex conjugation on the coefficients. For example, we have

$$\left(z_1^3 - z_1z_2 + i\right)^* = z_1^3 - z_2z_1 - i.$$

Let

$$\mathbb{C}\langle z_1,\dots,z_d\rangle_h := \{p \in \mathbb{C}\langle z_1,\dots,z_d\rangle \mid p^* = p\}$$

be the set of *Hermitian elements*. They form a real vectorspace, but not an algebra (in case $d \geq 2$). Note that Hermitian elements do not necessarily have real coefficients, and polynomials with real coefficients are not necessarily Hermitian.

Into a non-commutative polynomial $p \in \mathbb{C}\langle z_1,\dots,z_d\rangle$ we can plug in a d-tuple of elements from any complex algebra, and obtain an element from this algebra as the result. We will restrict ourselves to matrix algebras here, i.e. we take $(A_1,\dots,A_d) \in \mathrm{Mat}_s(\mathbb{C})^d$ for some $s \geq 1$ and obtain

$$p(A_1,\dots,A_d) \in \mathrm{Mat}_s(\mathbb{C}).$$

The need for an involution and Hermitian elements becomes clear when trying to capture real phenomena. Indeed, if $A_1,\dots,A_d \in \mathrm{Her}_s(\mathbb{C})$ are Hermitian matrices and $p \in \mathbb{C}\langle z_1,\dots,z_d\rangle_h$ is Hermitian as well, then so is the result:

$$p(A_1,\dots,A_d) \in \mathrm{Her}_s(\mathbb{C}).$$

A Hermitian matrix is *positive semidefinite*, if all of its Eigenvalues are nonnegative; we denote this by $\succcurlyeq 0$. This is the right notion of positivity in our setup, so if

$$p(A_1,\dots,A_d) \succcurlyeq 0$$

we say that p is nonnegative at the (non-commutative) point $(A_1,\dots,A_d) \in \mathrm{Her}_s(\mathbb{C})^d$.

It is obvious that every matrix $A \in \mathrm{Her}_s(\mathbb{C})$ that can be written as an *Hermitian square*

$$A = B^*B$$

for some $B \in \mathrm{Mat}_s(\mathbb{C})$ is positive semidefinite; in fact every positive semidefinite matrix is of that form. So if

$$p = \sum_{i=1}^{n} q_i^* q_i$$

for certain $q_1,\dots,q_n \in \mathbb{C}\langle z_1,\dots,z_d\rangle$, we obtain

$$p(A_1,\dots,A_d) = \sum_i q_i(A_1,\dots,A_d)^* q_i(A_1,\dots,A_d) \succcurlyeq 0$$

for any $(A_1,\dots,A_d) \in \mathrm{Her}_s(\mathbb{C})^d$. So the set of *sums of Hermitian squares*

$$\Sigma\mathbb{C}\langle z_1,\dots,z_d\rangle^2 := \left\{ \sum_{i=1}^{n} q_i^* q_i \mid n \in \mathbb{N}, q_i \in \mathbb{C}\langle z_1,\dots,z_d\rangle \right\} \subseteq \mathbb{C}\langle z_1,\dots,z_d\rangle_h$$

only contains polynomials that are positive semidefinite on each Hermitian matrix tuple. The first surprising result, a global Positivstellensatz and a non-commutative analogue of Hilbert's 17th Problem, is due to Helton:

Theorem 4 ([14])**.** *Let $p \in \mathbb{C}\langle z_1,\dots,z_d\rangle_h$ and assume*

$$p(A_1,\dots,A_d) \succcurlyeq 0$$

for all $(A_1,\dots,A_d) \in \mathrm{Her}_s(\mathbb{C})^d$ and all $s \geq 1$. Then

$$p \in \Sigma\mathbb{C}\langle z_1,\dots,z_d\rangle^2.$$

In contrast to the commutative result, no denominator is needed in Helton's theorem. However, the natural notion of positivity is much stronger here than in Hilbert's 17th Problem, where positivity is only assumed on matrices of size 1, instead of matrices of all sizes. Note however that also in Helton's result one can bound the matrix size, depending only on d and the degree of p.

Let us now define free basic closed semialgebraic sets. In analogy to the above described commutative setup, we define for $p_1,\dots,p_r \in \mathbb{C}\langle z_1,\dots,z_d\rangle_h$ and $s \geq 1$

$$W_s(p_1,\dots,p_r) := \left\{ (A_1,\dots,A_d) \in \mathrm{Her}_s(\mathbb{C})^d \mid p_i(A_1,\dots,A_d) \succcurlyeq 0, i = 1,\dots,r \right\}.$$

A guiding principle in non-commutative geometry is to not consider matrices of one size alone, but all sizes at once. We thus define the *free basic closed semialgebraic set* defined by $p_1,\dots,p_r$ as

$$\mathrm{F}W(p_1,\dots,p_r) := (W_s(p_1,\dots,p_r))_{s=1}^{\infty}.$$

There is no known Positivstellensatz for positivity on general free basic closed semialgebraic sets, however certain results in special cases. One of them deals with the *matrix cube*, see for example [1]:

Theorem 5. *Assume $p \in \mathbb{C}\langle z_1, \ldots, z_d \rangle_h$ is nonnegative on*

$$\mathrm{FW}\left(1 - z_1^2, \ldots, 1 - z_d^2\right).$$

Then there exists a representation

$$p = \sum_i q_i^* q_i + \sum_{i,j} q_{ij}^* \left(1 - z_j^2\right) q_{ij}$$

for certain $q_i, q_{ij} \in \mathbb{C}\langle z_1, \ldots, z_d \rangle$.

Another such Positivstellensatz is explained in the next section, and we also refer to [1] for more examples and unified proofs.

A notion of free semialgebraic sets beyond free basic closed semialgebraic sets has not been established in the literature so far. Boolean combination of basic closed sets will surely have to be allowed, but maybe that is not yet the best possible notion. This becomes clear when trying to prove a free projection theorem. For any $s \geq 1$ we apply the projection map

$$\begin{aligned} \pi_s \colon \mathrm{Her}_s(\mathbb{C})^d &\to \mathrm{Her}_s(\mathbb{C})^{d-1} \\ (A_1, \ldots, A_d) &\mapsto (A_1, \ldots, A_{d-1}) \end{aligned}$$

to $W_s(p_1, \ldots, p_r)$ and altogether obtain

$$\pi(\mathrm{FW}(p_1, \ldots, p_r)) := (\pi_s(W_s(p_1, \ldots, p_r)))_{s=1}^{\infty}.$$

How does such a projected set look like? Is it a Boolean combination of free basic closed sets? The answer to this question is no, and the whole topic does not look too encouraging. For example (see [6]), using free basic closed semialgebraic sets, Boolean combinations and projections, one can construct the set

$$(S_s)_{s=1}^{\infty}, \qquad S_s = \begin{cases} \mathrm{Her}_s(\mathbb{C}) & s \text{ prime} \\ \emptyset & \text{else.} \end{cases}$$

This set cannot be defined without projections, even if the language is enlarged by using trace, determinant and many other functions. Even more discouraging is the following result from [6]:

Theorem 6. *It is undecidable, whether a set constructed from free basic closed semialgebraic sets, Boolean combinations and projections is empty (at each level).*

A very recent positive result is [17]. Without going too much into the details, it states that quantifiers in non-commutative formulas can be eliminated, as long as the formula is evaluated at matrix tuples of *fixed size*. This is not a trivial result, since the variables in such formulas do not refer to the single matrix entries (where the result would follow from classical (commutative) quantifier elimination), but to matrices as a whole. However, the formula without quantifiers will depend on the matrix size. So the result does not imply a general (size independent) projection theorem.

A much more fruitful concept is *free convexity*, as we now explain in our last section.

3 Free Convexity

In this section we define the notion of a non-commutative convex set. Since definitions and results become cleaner for convex cones instead of convex sets, we restrict ourselves to cones here. As above we consider free sets

$$S = (S_s)_{s=1}^{\infty}, \quad S_s \subseteq \mathrm{Her}_s(\mathbb{C})^d \text{ for all } s \geq 1.$$

Matrix convexity of S is defined via two properties. A very reasonable assumption, even fulfilled for all free basic closed semialgebraic sets, is *closedness under direct sums*. For $\underline{A} = (A_1, \ldots, A_d) \in \mathrm{Her}_s(\mathbb{C})^d, \underline{B} = (B_1, \ldots, B_d) \in \mathrm{Her}_t(\mathbb{C})^d$ we define

$$\underline{A} \oplus \underline{B} := \left(\left(\begin{array}{c|c} A_1 & 0 \\ \hline 0 & B_1 \end{array} \right), \ldots, \left(\begin{array}{c|c} A_d & 0 \\ \hline 0 & B_d \end{array} \right) \right) \in \mathrm{Her}_{s+t}(\mathbb{C})^d.$$

S is closed under direct sums, if

$$\underline{A} \in S_s, \underline{B} \in S_t \;\Rightarrow\; \underline{A} \oplus \underline{B} \in S_{s+t}. \tag{C1}$$

The second condition resembles scaling with positive reals, but even connects the different levels of S. For $V \in \mathrm{Mat}_{s,t}(\mathbb{C})$ and $\underline{A} \in \mathrm{Her}_s(\mathbb{C})^d$ we define

$$V^* \underline{A} V := (V^* A_1 V, \ldots, V^* A_d V) \in \mathrm{Her}_t(\mathbb{C})^d.$$

The second condition then reads

$$\underline{A} \in S_s, V \in \mathrm{Mat}_{s,t}(\mathbb{C}) \;\Rightarrow\; V^* \underline{A} V \in S_t. \tag{C2}$$

If S fulfills (C1) and (C2), it is called a *matrix convex cone*. It is easily checked that each S_s is a classical convex cone in the real vector space $\mathrm{Her}_s(\mathbb{C})^d$ in this case. However, matrix convexity is a stronger condition in general, connecting

the different levels of S via (C2). Also note that a matrix convex cone is almost the same as an *abstract operator system* [8, 22], which only requires all S_s to be closed and salient with nonempty interior, additionally.

The most basic examples of matrix convex cones are *free spectrahedral cones* (or operator systems with a finite-dimensional realization, equivalently). For $M_1,\dots,M_d \in \mathrm{Her}_r(\mathbb{C})$ define

$$S_s(M_1,\dots,M_d) := \left\{(A_1,\dots,A_d) \in \mathrm{Her}_s(\mathbb{C})^d \mid M_1\otimes A_1+\cdots+M_d\otimes A_d \geqslant 0\right\}$$

and

$$\mathrm{F}S(M_1,\dots,M_d) = (S_s(M_1,\dots,M_d))_{s=1}^{\infty}.$$

Here, $\otimes$ denotes the Kronecker-/tensorproduct of matrices. The set $S_1(M_1,\dots,M_d)$ is known as a classical *spectrahedron*. Such sets are precisely the feasible sets of semidefinite programming. The free spectrahedron $\mathrm{F}S(M_1,\dots,M_d)$ is a non-commutative extension, precisely in the spirit as above. For free spectrahedra, there exists a nice Positivstellensatz. As in Theorems 4 and 5 above, we see that the natural notion of positivity in the non-commutative setup is strong enough to provide the best possible algebraic certificate (we do not cite the most general result and suppress some minor technical details for better readability):

Theorem 7 ([11])**.** *Let $M_1,\dots,M_d \in \mathrm{Her}_r(\mathbb{C})$ and $p \in \mathbb{C}\langle z_1,\dots,z_d\rangle_h$. If*

$$p(\underline{A}) \geqslant 0$$

for all $\underline{A} \in S_s(M_1,\dots,M_d)$ and all $s \geq 1$, in other words if p is nonnegative on the free spectrahedron $\mathrm{F}S(M_1,\dots,M_d)$, then there exists a representation

$$p = \sum_i q_i^* q_i + \sum_j f_j^* M f_j$$

where $q_i \in \mathbb{C}\langle z_1,\dots,z_d\rangle$, $f_j \in \mathbb{C}\langle z_1,\dots,z_d\rangle^r$ and

$$M := z_1M_1+\cdots+z_dM_d \in \mathrm{Her}_r\left(\mathbb{C}\langle z_1,\dots,z_d\rangle\right).$$

Sometimes facts about classical spectrahedra can only be learned by extending them to the non-commutative setup. One such instance is the *containment problem* for spectrahedra, a problem appearing in different areas of (applied) mathematics [16]. Given $M_1,\dots,M_d \in \mathrm{Her}_r(\mathbb{C})$ and $N_1,\dots,N_d \in \mathrm{Her}_t(\mathbb{C})$, how can one check efficiently whether

$$S_1(M_1,\dots,M_d) \subseteq S_1(N_1,\dots,N_d) \tag{A}$$

holds? Since spectrahedra are generalizations of polyhedra (which appear in the case of commuting coefficient matrices), this includes the problem of polyhedral

containment. An important algorithm to solve this problem was proposed in [2]. Instead of checking (A) once checks

$$\exists V_1,\dots,V_n \in \mathrm{Mat}_{r,t}(\mathbb{C}): \qquad \sum_j V_j^* M_i V_j = N_i \text{ for } i=1,\dots,d. \tag{B}$$

It is obvious that (B) implies (A). Condition (B) can be transformed into a semidefinite optimization problem, and thus often solved efficiently. It was however known that (A) and (B) are not equivalent, so the answer to (B) could be *no*, where the answer to (A) is *yes*. A much better understanding of the method was gained through the following result (again we suppress some minor technical details):

Theorem 8 ([9])**.** *Condition (B) is equivalent to*

$$\mathrm{FS}(M_1,\dots,M_d) \subseteq \mathrm{FS}(N_1,\dots,N_d). \tag{A'}$$

Inclusion is meant level-wise here, i.e. $S_s(M_1,\dots,M_d) \subseteq S_s(N_1,\dots,N_d)$ *for all* $s \geq 1$.

This result mostly relies on Choi's characterization of completely positive maps between matrix algebras [4]. The insight of Theorem 8 can now be used to determine instances in which (A) and (B) are equivalent nonetheless. For this let $C \subseteq \mathbb{R}^d$ be a convex cone. There is one smallest and one largest matrix convex set with C at level one. Indeed define

$$C_s^{\min} := \left\{ \sum_i c_i^t \otimes P_i \mid c_i \in C, P_i \in \mathrm{Her}_s(\mathbb{C}), P_i \geqslant 0 \right\}$$

and

$$C_s^{\max} := \left\{ \underline{A} \in \mathrm{Her}_s(\mathbb{C})^d \mid v^* \underline{A} v \in C \text{ for all } v \in \mathbb{C}^s \right\}.$$

Then

$$C^{\min} := \left(C_s^{\min}\right)_{s=1}^{\infty} \quad \text{and} \quad C^{\max} := (C_s^{\max})_{s=1}^{\infty}$$

are easily checked to be the smallest/largest such matrix convex set. Now assume

$$C = S_1(M_1,\dots,M_d) \subseteq \mathbb{R}^d$$

is a (classical) spectrahedral cone with $\mathrm{FS}(M_1,\dots,M_d) = C^{\min}$. In that case, condition (A) implies

$$\mathrm{FS}(M_1,\dots,M_d) = C^{\min} \subseteq \mathrm{FS}(N_1,\dots,N_d)$$

and thus (B), by Theorem 8. On the other hand, if

$$C^{\min} \subsetneq \mathrm{FS}(M_1,\dots,M_d)$$

it can be shown by the non-commutative separation theorem from [7], that there exist matrices $N_1,\dots,N_d$ with

$$C = S_1(M_1,\dots,M_d) = S_1(N_1,\dots,N_d)$$

and

$$\mathrm{FS}(M_1,\dots,M_d) \not\subseteq \mathrm{FS}(N_1,\dots,N_d).$$

In such an instance the answer to (B) is no, whereas the answer to (A) is yes. So the method from [2] works reliably if any only if $\mathrm{FS}(M_1,\dots,M_d)$ is the smallest matrix convex cone over the classical spectrahedron $S_1(M_1,\dots,M_d)$. Unfortunately, this happens very rarely, already for polyhedral cones:

Theorem 9 ([8])**.** *Assume $C = S_1(M_1,\dots,M_d) \subseteq \mathbb{R}^d$ is polyhedral. Then*

$$\mathrm{FS}(M_1,\dots,M_d) = C^{\min},$$

if and only if C is a simplex cone, i.e. has only d extremal rays.

The last theorem also has some surprising application in theoretical quantum physics. The state of a bipartite quantum system is usually described by a positive semidefinite matrix

$$\rho \in \mathrm{Mat}_r(\mathbb{C}) \otimes \mathrm{Mat}_s(\mathbb{C}) \cong \mathrm{Mat}_{rs}(\mathbb{C}).$$

So ρ can be written as

$$0 \leqslant \rho = \sum_{i=1}^{n} M_i \otimes A_i$$

with $M_i \in \mathrm{Mat}_r(\mathbb{C}), A_i \in \mathrm{Mat}_s(\mathbb{C})$. Although ρ is supposed to be positive semidefinite and in particular Hermitian, this is not necessarily true for the M_i, A_i. If there exists a different such representation, where all the M_i, A_i are positive semidefinite as well, then ρ is called *separable*, otherwise it is *entangled*. The smallest possible n in the representation of ρ above is called the *tensor rank* of ρ. A corollary of Theorem 9 now reads as follows:

Theorem 10 ([5])**.** *Every bipartite quantum state of tensor rank* 2 *is separable.*

In fact

$$0 \leqslant \rho = M_1 \otimes A_1 + M_2 \otimes A_2$$

just means that $(A_1, A_2) \in S_s(M_1, M_2)$. Now since the convex cone

$$C = S_1(M_1, M_2) \subseteq \mathbb{R}^2$$

is automatically a simplex cone, we obtain $(A_1, A_2) \in C_s^{\min}$ from Theorem 9. Writing down a representation in this smallest matrix convex cone and using bilinearity of the tensor product immediately implies the result.

Let us close with a result about non-commutative polytopes and polyhedra. The theorem of Minkowski-Weyl (see for example [26]) states that every polyhedral cone $C \subseteq \mathbb{R}^d$ is finitely generated, and vice versa. In other words, the notions *polyhedral* and *polytopal* coincide for convex cones. Now a short contemplation reveals that $C^{\min}$ is a good generalization of the notion *polytope/finitely generated* to the non-commutative setup, whereas $C^{\max}$ corresponds to the *polyhedral* notion. Interestingly, these two notions differ almost always, already at the first level of non-commutativity:

Theorem 11 ([8, 21, 15])**.** *Let $C \subseteq \mathbb{R}^d$ be a convex cone.*
(i) If C is a simplex cone, then $C^{\min} = C^{\max}$. Otherwise $C^{\min} \neq C^{\max}$.
(ii) If C is polyhedral but not a simplex cone, then $C_2^{\min} \subsetneq C_2^{\max}$.

As a concluding remark, we note that the methods used in the non-commutative setup differ quite strongly from the ones in the commutative theory. Many of the results are proven by functional-analytic methods, such as GNS-constructions, dilations, and the theory of completely positive maps and operator algebras. Sometimes results and examples from group theory and the theory of C^*-algebras can be useful. All in all, the whole area is not yet mature, many interesting results and methods are hopefully developed in the coming years.

References

[1] V. Alekseev, T. Netzer, T., A.Thom, Quadratic modules, C^*-algebras and free convexity, *Trans. Amer. Math. Soc.*, to appear.

[2] A. Ben-Tal, A. Nemirovski, On tractable approximations of uncertain linear matrix inequalities affected by interval uncertainty, *SIAM J. Optim.* **12** (2002), 811–833.

[3] J. Bochnak, M. Coste, M.-F. Roy, *Real algebraic geometry*, Ergebnisse der Mathematik und ihrer Grenzgebiete (3) **36**, Springer-Verlag, Berlin, 1998.

[4] M. D. Choi, Completely positive linear maps on complex matrices, *Linear Algebra and Appl.* **10** (1975), 285–290.

[5] G. De Las Cuevas, T. Drescher, T. Netzer, Matrix product density operators of bond dimension two are separable, in preparation.

[6] T. Drescher, T. Netzer, A. Thom, On projections of free semialgebraic sets, preprint, 2017.

[7] E. G. Effros, S. Winkler, Matrix convexity: operator analogues of the bipolar and Hahn-Banach theorems, *J. Funct. Anal.* **144** (1997), 117–152.

[8] T. Fritz, T. Netzer, A. Thom, Spectrahedral containment and operator systems with finite-dimensional realization, *SIAM J. Appl. Algebra Geom.* **1** (2017), 556–574.

[9] J. W. Helton, I. Klep, S. McCullough, The matricial relaxation of a linear matrix inequality, *Math. Program.* **138** (2013), 401–445.

[10] J. W. Helton, I. Klep, S. McCullough, Free convex algebraic geometry, *Semidefinite optimization and convex algebraic geometry*, *MOS-SIAM Ser. Optim.* 13, 341–405, SIAM, Philadelphia, PA, 2013.

[11] J. W. Helton, I. Klep, S. McCullough, The convex Positivstellensatz in a free algebra, *Adv. Math.* **231** (2012), 516–534.
[12] J. W. Helton, I. Klep, S. McCullough, Every convex free basic semi-algebraic set has an LMI representation, *Ann. of Math. (2)* **176** (2012), 979–1013.
[13] J. W. Helton, S. McCullough, M. Putinar, A non-commutative Positivstellensatz on isometries, *J. Reine Angew. Math.* **568** (2004), 71–80.
[14] J. W. Helton, "Positive" noncommutative polynomials are sums of squares, Ann. of Math. (2) **156** (2002), 675–695.
[15] B. Huber, T. Netzer, A note on non-commutative polytopes and polyhedra, preprint, 2018.
[16] K. Kellner, Th. Theobald, Ch. Trabandt, Containment problems for polytopes and spectrahedra, *SIAM J. Optim.* **23** (2013), 1000–1020.
[17] I. Klep, M. Tressl, A model theoretic perspective of matrix rings, preprint, 2018.
[18] M. Marshall, *Positive polynomials and sums of squares*, Mathematical Surveys and Monographs **146**, American Mathematical Society, Providence, RI, 2008
[19] T. Netzer, Reelle Algebraische Geometrie, Lecture Notes, https://algebra-mathematics.uibk.ac.at/images/documents/teaching/tim_netzer/RAG.pdf.
[20] T. Netzer, Real algebraic geometry and its applications, *Internat. Math. Nachrichten* **232** (2016), 1–18.
[21] B. Passer, O. M. Shalit, B. Solel, Minimal and maximal matrix convex sets, *J. Funct. Anal.* **274** (2018), 3197–3253.
[22] V. Paulsen, *Completely bounded maps and operator algebras*, Cambridge Studies in Advanced Mathematics **78**, Cambridge University Press, Cambridge, 2002.
[23] A. Prestel, Ch. N. Delzell, *Positive polynomials*, Springer Monographs in Mathematics, From Hilbert's 17th problem to real algebra, Springer-Verlag, Berlin, 2001.
[24] M. Ramana, A. J. Goldman, Some geometric results in semidefinite programming, *J. Global Optim.* **7** (1995), 33–50.
[25] K. Schmüdgen, Noncommutative real algebraic geometry–some basic concepts and first ideas, Emerging applications of algebraic geometry, IMA Vol. Math. Appl. **149**, Springer, New York, 2009, 325–350.
[26] A. Schrijver, *Theory of linear and integer programming*, Wiley-Interscience Series in Discrete Mathematics, John Wiley & Sons, Ltd., Chichester, 1986.
[27] Handbook of semidefinite programming, International Series in Operations Research & Management Science, 27, Wolkowicz, H., Saigal, R,, Vandenberghe, L. (Eds.), Kluwer Academic Publishers, Boston, MA, 2000.

Author's address:
University of Innsbruck
Department of Mathematics
Technikerstr. 13
A-6020 Innsbruck
email tim.netzer@uibk.ac.at

Buchbesprechungen

X. Hou: Lectures on Finite Fields. (Graduate Studies in Mathematics, Vol. 190.) American Mathematical Society, Providence (USA), 2018, 240 S. ISBN 978-1-4704-4289-7 P/b $ 83.

The theory of Finite Fields and Their Applications is a research area of increasing interest and deep impact to areas such as coding theory, cryptography, and quasi-Monte Carlo methods.

The book under review is based on a graduate course taught by the author and complements the encyclopedic books of Lidl and Niederreiter, Shparlinski and the Handbook of Finite Fields edited by Mullen and Panario.

Chapter 1 describes the basic structural properties of finite fields. Chapter 2 is entitled polynomials over finite fields and deals with the number of irreducible polynomials, factorization of polynomials, permutation polynomials and Payne's theorem. Chapter 3 studies Gauss sums.

The remaining chapters are not standard for a textbook on finite fields and make this book particularly valuable. Chapter 4 is an introduction to algebraic number theory, Chapter 5 proves the Ax-Katz-Theorem on the number of zeros of polynomials over finite fields, and Chapter 6 is an introduction to classical groups.

A. Winterhof (Linz)

T. Bühler, D. A. Salamon: Functional Analysis. (Graduate Studies in Mathematics, Vol. 191.) American Mathematical Society, Providence (USA), 2018, 466. S. ISBN 978-1-4704-4190-6 P/b $ 83.

This monograph is a very readable exposition of linear functional analysis which covers all the standard topics one would like to discuss in an intensive first course on the topic. The book starts with topological foundations and metric spaces and then discusses the classical three principles of linear functional analysis (uniform boundedness, open mapping, Hahn-Banach). Next, there is a chapter on weak topologies and the Banach-Alaoglu theorem. The exposition continues with a basic introduction to Fredholm theory. The second half of the book starts with spectral theory for bounded operators and in particular treats the Dunford integral (holomorphic functional calculus) and the continuous functional calculus for self-adjoint operators. As is perhaps appropriate for an introductory text on the subject, only in the second to last chapter unbounded operators are mentioned for the first time (with a focus on the self-adjoint case). The main result here is the spectral theorem for (unbounded) self-adjoint operators which is approached via projection-valued measures. Finally, the last chapter gives a basic introduction to semigroup theory. The book is very well written with detailed proofs, a lot of examples, and interesting problems.

R. Donninger (Wien)

T.-P. Tsai: Lectures on the Navier-Stokes Equations. (Graduate Studies in Mathematics, Vol. 192.) American Mathematical Society, Providence (USA), 2018, 224. S. ISBN 978-1-4704-3096-2 P/b $ 83.

The book gives an up to date introduction to some of the most interesting aspects of the mathematical theory of the Navier-Stokes equation. It grew out of lecture notes and is therefore ideal for a graduate course for students who have been exposed to some PDE theory. At Austrian universities the most appropriate setting would be a topics course on partial differential equations. The book does not strive for an encyclopedic treatment (which would be quite unreadable) but on the contrary focuses on some aspects of current interest and goes straight to the point. To keep the exposition at a reasonable length, many preliminaries are just mentioned and for the proofs the reader is referred to other sources. This is absolutely necessary so as not to get lost in preparatory theory without ever getting to the point.

After a quick introduction to the physical origins of the Navier-Stokes equation and its basic properties, the author treats the different solution concepts of weak, strong, and mild solutions. After less than a hundred pages the reader arrives at one of the most challenging problems of contemporary mathematics: Navier-Stokes regularity. The book discusses the state of the art on this issue and describes a number of current results on partial regularity. From this point onwards the book stays at the forefront of current research and discusses bifurcation, self-similar solutions, uniqueness of weak solutions, and axisymmetric problems. The style of writing is very lively and one feels the excitement in this highly active field. In summary, the book is a very accessible introduction to some key aspects of this exciting area of contemporary research in partial differential equations.

R. Donninger (Wien)

M. Lorenz: A Tour of Representation Theory. (Graduate Studies in Mathematics, Vol. 193.) American Mathematical Society, Providence (USA), 2018, 654. S. ISBN 978-1-4704-3680-3 P/b $ 94.

The book of Lorenz on the theory of representations has the aim to give a panoramic viewpoint of this fascinating topic. So the author discusses representations of Algebras, Groups, Lie algebras and Hopf algebras in separate parts of the book.

The first part is devoted to the representation of Algebras. Beside basic topics such as irreducibility, semisimplicity and characters also the representation of projective modules and so-called Frobenius algebras is discussed including Wedderburns structure theorem.

Part two deals mostly with the finite dimensional representation of finite groups and the material covered in chapter 3 can be found in most other books on this topic. In the fourth chapter the author presents a rather new approach due to

Okounkov and Vershik and proves classical results on the representations of symmetric groups.

The topic of the third part are representations of Lie algebras. After an introduction to Lie algebras (no prior knowledge of Lie algebras is expected) and basic facts about their representations, Lorenz turns to the semisimple Lie algebras, root systems and the representation of semisimple Lie algebras.

In the final chapter the theory of Hopf algebras and their representations is presented. In particular, the author applies the theory to affine group schemes and gives a proof of Chevally's theorem that the tensor product of two finite-dimensional completely reducible representations of a group over a field of characteristic zero is again completely reducible.

The book is very well written and explains everything in great detail and illustrates the theory with many examples. Also numerous exercises help to better understand the material. Although the book gives a rather complete overview of the algebraic side of representations, the analytic side of representation theory, such as the representation theory of C^* algebras, is absent.

V. Ziegler (Salzburg)

G. T. Lee: Abstract Algebra. An Introductory Course. (Springer Undergraduate Mathematics Series.) Springer International Publishing, Cham, 2018, 301 S. ISBN 978-3-319-77648-4 P/b € 38,49.

This book aims for beginners in abstract algebra. One of its features is that it does not assume any previous knowledge of the subject. It starts in Part I with a reminder of relations, functions, integers, and modular arithmetic. Part II gives an introduction to group theory including the classification of finite abelian groups, the simpleness of the alternating group, and Sylow theory. Part III continues with rings including a chapter on special domains (polynomial rings, euclidean rings, UFDs, and PIDs). Part IV covers some basic field theory. This part is somewhat non standard since it starts with irreducibility of polynomials (over $\mathbb{Q}$, $\mathbb{C}$, and $\mathbb{F}_q$). Moreover, some linear algebra is presented in order to keep the book self-contained. In the last part public key cryptosystems and straightedge-and-compass constructions are covered. The complex numbers and matrix algebras are introduced in appendices, again, to keep the book self-contained.

The book is very clearly written. The author successfully presents the material in an appealing way. A big number of examples enriches the text and enlightens the key topics. Exercises of different level are included at the end of each chapter and solutions to approximately half of the exercises are included at the very end of the book. In summary, by all means the book can definitely be recommended as text book for a first introduction to abstract algebra.

C. Fuchs (Salzburg)

M. Friedman: A History of Folding in Mathematics. Mathematizing the Margins. Birkhäuser, Basel, 2018, 419 S. ISBN 978-3-319-72486-7 H/b € 155,14.

Die wissenschaftliche Auseinandersetzung mit dem Papierfalten ist in den letzten 20 Jahren förmlich explodiert, und so ist es nicht überraschend, dass man sich dem Versuch widmet, eine einigermaßen vollständige Historie dieses Spezialbereichs zu schreiben. Das vorliegende Werk bietet auch einen recht ausführlichen Überblick über die verschiedenen Aspekte dieses Themas, von den ersten Schritten in diese Richtung vor über einem Jahrhundert bis hin zum aktuellen Stand der Forschung. Es werden dabei Einblicke in didaktische Überlegungen, Fragen der Axiomatik und Konstruierbarkeit und vieles mehr geboten. In diesem Buch findet jede am Fach interessierte Person eine große Fülle an Information und auch eine Vielzahl an Hinweisen zu weiteren Informationsquellen. Wenn eine persönliche Bemerkung an dieser Stelle erlaubt sein darf, bin ich als Rezensent etwas darüber enttäuscht, dass meine eigenen Beiträge zu diesem Thema keine Beachtung gefunden haben, aber eine Lektüre ist sicherlich uneingeschänkt zu empfehlen.

R. Geretschäger (Graz)

A. George, D. J. Vellemann: Zur Philosophie der Mathematik. Logizismus, Intuitionismus, Finitismus, Gödel'sche Unvollständigkeitssätze. Springer Spektrum, Berlin/Heidelberg, 2018, 210 S. ISBN 978-3-662-56236-9 P/b € 29,99.

Lässt der Haupttitel des Buchs noch einigen Interpretationsspielraum, was man *Zur Philosophie der Mathematik* zu erwarten habe, so umreißt das der Untertitel sehr präzise. Denn vier der (abgesehen von Einleitung und Schluss) sechs Hauptkapitel des Buchs tragen die Titel *Logizismus, Intuitionismus, Finitismus* und *Die Unvollständigkeitssätze*. Dazwischen eingeschoben sind die beiden verbleibenden Hauptkapitel *Mengenlehre* und *Intuitionistische Mathematik*.

Wie die Autoren im Vorwort selbst betonen, geht es ihnen weniger um eine Darstellung der historischen Entwicklungen als um einen möglichst direkten Zugang sowohl zu den philosophischen Anliegen als auch zu den mathematischen Ideen, die mit der großen Revision der Grundlagen der Mathematik in den Jahrzehnten rund um das Jahr 1900 dominierten und die besonders eng mit den Namen Frege, Brouwer und Hilbert verknüpft sind. Die folgenden Bemerkungen mögen die Inhalte etwas genauer beschreiben.

Im Logizismus- sowie dem darauffolgenden Mengenlehre-Kapitel wird deutlich, wie sehr Frege in seiner berühmten *Begriffsschrift* aus dem Jahr 1879 schon substanzielle Ideen vorweggenommen hatte, die sich erst später in ihrer mengentheoretischen Ausformulierung durch seine Nachfolger durchsetzten. Er selbst scheiterte bei seinen Versuchen, die Arithmetik aus der reinen Logik zu begründen, im Wesentlichen an der Russellschen Antinomie. Auch die beiden Kapitel über die philosophischen Anliegen des Intuitionismus und über die Besonderheiten der daraus resultierenden Mathematik vermitteln ein sehr klares Bild von ihrem Ge-

genstand. Im Kapitel über den Finitismus liegt, der Überschrift entsprechend, der Schwerpunkt auf Hilberts Bestrebungen, möglichst die gesamte Mathematik mit finitären Mitteln zu rechtfertigen, und weniger auf der axiomatischen Methode, die über das Schlagwort *Formalismus* gleichfalls stark mit Hilberts epochalem Wirken assoziiert wird und die vielleicht als der eigentliche Sieger aus den Scharmützeln zwischen den drei Parteien hervorging. Der in diesem Zusammenhang so wichtige Gödelsche Vollständigkeitssatz findet nur nebenbei im letzten Hauptkapitel Erwähnung. Dieses ist den Unvollständigkeitssätzen samt Beweisen gewidmet. Das Schlusskapitel ist eher philosophisch ausgerichtet, zeigt manche Ausblicke auf, lässt aber bewusst auch viele Fragen offen.

Das Buch setzt wenig Vorwissen voraus und spricht entsprechend eine breite Leserschaft an, die über das fachmathematische Publikum deutlich hinausreicht.

R. Winkler (Wien)

Nachrichten der Österreichischen Mathematischen Gesellschaft

Gilbert Helmberg 1928–2019

Am 18. Februar 2019 ist Gilbert Helmberg, em.o.Univ.-Prof. der Leopold-Franzens-Universität Innsbruck, im 90. Lebensjahr verstorben. Neben seinen vielfältigen akademischen Tätigkeiten war er von 1994-1997 Vorsitzender der ÖMG, deren Mitglied er seit 1952 war. Für das nächste Heft der IMN ist ein ausführlicher Nachruf geplant.

Walter Knödel 1926–2018

Am 19. Oktober 2018 verstarb Prof. Dr. Walter Knödel, der nach Studium, Promotion (bei Edmund Hlawka) und Habilitation an der Univerität Wien im Jahr 1961 auf den Lehrstuhl für Instrumentelle Mathematik an der Technischen Hochschule Stuttgart berufen wurde und unter anderem Gründungsdekan der Fakultät für Informatik der Universität Stuttgart war.

Neue Mitglieder

Pucher Dunja – Wulfengasse 13/7, 9020 Klagenfurt. geb. 1980. Diplomstudium der Betriebswirtschaft an der Universität Zagreb von 1998 bis 2005. Derzeit Bachelorstudium Mathematik an der Universität Klagenfurt. email *dupucher@edu.aau.at*

Eichmair Michael, Univ.-Prof. Dr. – c/o Oskar-Morgenstern-Platz 1, 1090 Wien. geb. 1983. Studium der Mathematik in Linz und London von 2000 bis 2003. Doktorat an der Stanford University 2008. C.L.E. Moore Instructor am MIT, Assistenzprofessor an der ETH Zürich, seit 2015 Professor für Globale Analysis und Differentialgeometrie an der Universität Wien. START-Preis 2016, Förderungspreis der ÖMG 2017. Für weitere Details siehe https://medienportal.univie.ac.at/uniview/professuren/cv/artikel/univ-prof-dr-michael-eichmair/. email *michael.eichmair@univie.ac.at*

van den Berg Robert, Dr. – Hauptstraße, 2351 Wiener Neudorf. geb. 1971. Bachelor und Master in Angewandter Mathematik & Statistik, promovierter Leistungsphysiologe. email *robertvandenberg71@gmail.com*

Izmestiev Ivan, Dr. – c/o Chemin du Musée 23, CH-1700 Fribourg. geb. 1975. Doktorat an der Moscow State University 2001. Verschiedene Stellen in Deutschland und in der Schweiz. Derzeit Oberassistent an der Universität Freiburg in der Schweiz. email *ivan.izmestiev@gmail.com*