1

Introduction

1.1 The Point of Departure

In one of the papers initiating the study of the continuous wavelet transform on the real line, Grossmann, Morlet and Paul [60] considered systems $(\psi_{b,a})_{b,a\in\mathbb{R}\times\mathbb{R}'}$ arising from a single function $\psi \in \mathrm{L}^2(\mathbb{R})$ via

$$\psi_{b,a}(x) = |a|^{-1/2}\psi\left(\frac{x-b}{a}\right) \quad .$$

They showed that every function ψ fulfilling the **admissibility condition**

$$\int_{\mathbb{R}'} \frac{|\widehat{\psi}(\omega)|^2}{|\omega|} d\omega = 1 \quad , \tag{1.1}$$

where $\mathbb{R}' = \mathbb{R} \setminus \{0\}$, gives rise to an **inversion formula**

$$f = \int_{\mathbb{R}} \int_{\mathbb{R}'} \langle f, \psi_{b,a}\rangle \psi_{b,a} \frac{da}{|a|^2} db \quad , \tag{1.2}$$

to be read in the weak sense. An equivalent formulation of this fact is that the **wavelet transform**

$$f \mapsto V_\psi f \quad , \quad V_\psi f(b,a) = \langle f, \psi_{b,a}\rangle$$

is an isometry $\mathrm{L}^2(\mathbb{R}) \to \mathrm{L}^2(\mathbb{R}\times\mathbb{R}', db\frac{da}{|a|^2})$. As a matter of fact, the inversion formula was already known to Calderón [27], and its proof is a more or less elementary exercise in Fourier analysis.

However, the admissibility condition as well as the choice of the measure used in the reconstruction appear to be somewhat obscure until read in group-theoretic terms. The relation to groups was pointed out in [60] –and in fact earlier in [16]–, where it was noted that $\psi_{b,a} = \pi(b,a)\psi$, for a certain representation π of the affine group G of the real line. Moreover, (1.1) and (1.2) have natural group-theoretic interpretations as well. For instance, the measure used for reconstruction is just the left Haar measure on G.

Hence, the wavelet transform is seen to be a special instance of the following construction: Given a (strongly continuous, unitary) representation $(\pi, \mathcal{H}_\pi)$ of a locally compact group G and a vector $\eta \in \mathcal{H}_\pi$, we define the **coefficient operator**

$$V_\eta : \mathcal{H}_\pi \ni \varphi \mapsto V_\eta \varphi \in C_b(G) \quad , \quad V_\eta \varphi(x) = \langle \varphi, \pi(x)\eta \rangle \quad .$$

Here $C_b(G)$ denotes the space of bounded continuous functions on G.

We are however mainly interested in inversion formulae, hence we consider V_η as an operator $\mathcal{H}_\pi \to \mathrm{L}^2(G)$, with the obvious domain $\mathrm{dom}(V_\eta) = \{\varphi \in \mathcal{H}_\pi : V_\eta \varphi \in \mathrm{L}^2(G)\}$. We call η **admissible** whenever $V_\eta : \mathcal{H} \to \mathrm{L}^2(G)$ is an isometric embedding, and in this case V_η is called **(generalized) wavelet transform**. While the definition itself is rather simple, the problem of identifying admissible vectors is highly nontrivial, and the question whether these vectors exist for a given representation does not have a simple general answer. It is the main purpose of this book to develop in a systematical fashion criteria to deal with both problems.

As pointed out in [60], the construction principle for wavelet transforms had also been studied in mathematical physics, where admissible vectors η are called **fiducial vectors**, systems of the type $\{\pi(x)\eta : x \in G\}$ **coherent state systems**, and the corresponding inversion formulae **resolutions of the identity**; see [1, 73] for more details and references.

Here the earliest and most prominent examples were the original coherent states obtained by time-frequency shifts of the Gaussian, which were studied in quantum optics [114]. Perelomov [97] discussed the existence of resolutions of the identity in more generality, restricting attention to irreducible representations of unimodular groups. In this setting **discrete series** representations, i.e., irreducible subrepresentations of the regular representation λ_G of G turned out to be the right choice. Here every nonzero vector is admissible up to normalization. Moreover, Perelomov devised a construction which gives rise to resolutions of the identity for a large class of irreducible representations which were not in the discrete series. The idea behind this construction was to replace the group as integration domain by a well-chosen quotient, i.e., to construct isometries $\mathcal{H}_\pi \hookrightarrow \mathrm{L}^2(G/H)$ for a suitable closed subgroup H. In all of these constructions, irreducibility was essential: Only the well-definedness and a suitable intertwining property needed to be proved, and Schur's lemma would provide for the isometry property.

While we already remarked that [60] was not the first source to comment on the role of the affine group in constructing inversion formulae, suitably general criteria for nonunimodular groups were missing up to this point. Grossmann, Morlet and Paul showed how to use the orthogonality relations, established for these groups by Duflo and Moore [38], for the characterization of admissible vectors. More precisely, Duflo and Moore proved the existence of a uniquely defined unbounded selfadjoint operator C_π associated to a discrete series representation such that a vector η is admissible iff it is contained in the domain of C_π, with $\|C_\pi \eta\| = 1$. A second look at the admissibility condition (1.1) shows that in the case of the wavelet transform on $\mathrm{L}^2(\mathbb{R})$ this operator is given on the Plancherel transform side by multiplication with $|\omega|^{-1/2}$. This framework allowed to construct analogous transforms in a variety of settings, which was to become an active area of research in the subsequent years; a by no means complete list of references is [93, 22, 25, 48, 68, 49, 50, 51, 83, 7, 8]. See also [1] and the references therein.

However, it soon became apparent that admissible vectors exist outside the discrete series setting. In 1992, Mallat and Zhong [92] constructed a transform related to the original continuous wavelet transform, called the **dyadic wavelet transform**. Starting from a function $\psi \in \mathrm{L}^2(\mathbb{R})$ satisfying the **dyadic admissibility condition**

$$\sum_{n\in\mathbb{Z}} |\widehat{\psi}(2^n\omega)|^2 = 1 \quad , \quad \text{for almost every } \omega \in \mathbb{R} \tag{1.3}$$

one obtains the (weak-sense) inversion formula

$$f = \int_{\mathbb{R}} \sum_{n\in\mathbb{Z}} \langle f, \psi_{b,2^n} \rangle \psi_{b,2^n} 2^{-n} db \quad , \tag{1.4}$$

or equivalently, an isometric **dyadic wavelet transform** $\mathrm{L}^2(\mathbb{R}) \to \mathrm{L}^2(\mathbb{R} \times \mathbb{Z}, db2^{-n}dn)$, where dn denotes counting measure. Clearly the representation behind this transform is just the restriction of the above representation π to the closed subgroup $H = \{(b, 2^n) : b \in \mathbb{R}, n \in \mathbb{Z}\}$ of G, and the measure underlying the dyadic inversion formula is the left Haar measure of that subgroup. However, in one respect the new transform is fundamentally different: The restriction of π to H is no longer irreducible, in fact, it does not even contain irreducible subrepresentations (see Example 2.36 for details). Therefore (1.3) and (1.4), for all the apparent similarity to (1.1) and (1.2), cannot be treated in the same discrete series framework.

The example by Mallat and Zhong, together with results due to Klauder, Isham and Streater [67, 74], was the starting point for the work presented in this book. In each of these papers, a more or less straightforward construction led to admissibility conditions – similar to (1.1) and (1.3) – for representations which could not be dealt with by means of the usual discrete series arguments. The initial motivation was to understand these examples under a representation-theoretic perspective, with a view to providing a general strategy for the systematic construction of wavelet transforms.

The book departs from a few basic realizations: Any wavelet transform V_η is a unitary equivalence between π and a subrepresentation of λ_G, the left regular representation of G on $\mathrm{L}^2(G)$. Hence, the **Plancherel decomposition** of the latter into a direct integral of irreducible representations should play a central role in the study of admissible vectors, as it allows to analyze invariant subspaces and intertwining operators.

A first hint towards direct integrals had been given by the representations in [67, 74], which were constructed as direct integrals of irreducible representations. However, the particular choice of the underlying measure was not motivated, and it was unclear to what extent these constructions and the associated admissibility conditions could be generalized to other groups. Properly read, the paper by Carey [29] on reproducing kernel subspaces of $\mathrm{L}^2(G)$ can be seen as a first source discussing the role of Plancherel measure in this context.

1.2 Overview of the Book

The contents of the remaining chapters may be roughly summarized as follows:

2. Introduction to the group-theoretic approach to the construction of continuous wavelet transforms. Embedding the discussion into $\mathrm{L}^2(G)$. Formulation of a list of tasks to be solved for general groups. Solution of these problems for the toy example $G = \mathbb{R}$.
3. Introduction to the Plancherel transform for type I groups, and to the necessary representation-theoretic machinery.
4. Plancherel inversion and admissibility conditions for type I groups. Existence and characterization of admissible vectors for this setting.
5. Examples of admissibility conditions in concrete settings, in particular for quasiregular representations.
6. Sampling theory on the Heisenberg group.

Chapter 2 is concerned with the collection of basic notions and results, concerning coefficient operators, inversion formulae and their relation to convolution and the regular representations. In this chapter we formulate the problems which we intend to address (with varying degrees of generality) in the subsequent chapters. We consider existence and characterization of inversion formulae, the associated reproducing kernel subspaces of $\mathrm{L}^2(G)$ and their properties, and the connection to discretization of the continuous transforms and sampling theorems on the group. Support properties of the arising coefficient functions are also an issue. Section 2.7 is crucial for the following parts: It discusses the solution of the previously formulated list of problems for the special case $G = \mathbb{R}$. It turns out that the questions mostly translate to elementary problems in real Fourier analysis.

Chapter 3 provides the "Fourier transform side" for locally compact groups of type I. The Fourier transform of such groups is obtained by integrating functions against irreducible representations. The challenge for Plancherel theory is to construct from this a unitary operator from $\mathrm{L}^2(G)$ onto a suitable direct integral space. This problem may be seen as analogous to the case of the reals, where the tasks consists in showing that the Fourier transform defined on $\mathrm{L}^1(\mathbb{R})$ induces a unitary operator $\mathrm{L}^2(\mathbb{R}) \to \mathrm{L}^2(\mathbb{R})$. However, for arbitrary locally compact groups the right hand side first needs to be constructed, which involves a fair amount of technique. The exposition starts from a representation-theoretic discussion of the toy example, and during the exposition to follow we refer repeatedly to this initial example.

Chapter 4 contains a complete solution of the existence and characterization of admissible vectors, at least for type I groups and up to unitary equivalence. The technique is a suitable adaptation of the Fourier arguments used for the toy example. It relies on a pointwise Plancherel inversion formula, which in this generality has not been previously established. In the course of argument we derive new results concerning the Fourier algebra and Fourier inversion on type I locally compact groups, as well as an L^2-version of the convolution theorem, which allows a precise description of L^2-convolution operators, including domains, on the Plancherel transform side 4.18. We comment on an interpretation of the support properties obtained in Chapter 2 in connection with the so-called "qualitative uncertainty principle". Using existence and uniqueness properties of direct integral decompositions, we then describe a general procedure how to establish the existence and criteria for admissible vectors (Remark 4.30). We also show that these criteria in effect characterize the Plancherel measure, at least for unimodular groups. Section 4.5 shows how the Plancherel transform view allows a unified treatment of wavelet and Wigner transforms associated to nilpotent Lie groups.

Chapter 5 shows how to put the representation-theoretic machinery developed in the previous chapters to work on a much-studied class of concrete representations, thereby considerably generalizing the existing results and providing additional theoretic background. We discuss semidirect products of the type $\mathbb{R}^k \rtimes H$, with suitable matrix groups H. These constructions have received considerable attention in the past. However, the representation-theoretic results derived in the previous chapters allow to study generalizations, e.g. groups of the sort $N \rtimes H$, where N is a homogeneous Lie group and H is a one-parameter group of dilations on N. The discussion of the Zak-transform in the context of Weyl-Heisenberg frames gives further evidence for the scope of the general representation-theoretic approach.

The final chapter contains a discussion of sampling theorems on the Heisenberg group $\mathbb{H}$. We obtain a complete characterization of the closed leftinvariant subspaces of $\mathrm{L}^2(\mathbb{H})$ possessing a sampling expansion with respect to a lattice. Crucial tools for the proof of these results are provided by the theory of Weyl-Heisenberg frames.

1.3 Preliminaries

In this section we recall the basic notions of representation theory, as far as they are needed in the following chapter. For results from representation theory, the books by Folland [45] and Dixmier [35] will serve as standard references.

The most important standing assumptions are that *all locally compact groups in this book are assumed to be Hausdorff and second countable* and *all Hilbert spaces in this book are assumed to be separable.*

Hilbert Spaces and Operators

Given a Hilbert space $\mathcal{H}$, the space of bounded operators on it is denoted by $\mathcal{B}(\mathcal{H})$, and the operator norm by $\|\cdot\|_\infty$. $\mathcal{U}(\mathcal{H})$ denotes the group of unitary operators on $\mathcal{H}$. Besides the norm topology, there exist several topologies of interest on $\mathcal{B}(\mathcal{H})$. Here we mention the **strong operator topology** as the coarsest topology making all mappings of the form

$$\mathcal{B}(\mathcal{H}) \ni T \mapsto T\eta \in \mathcal{H} \quad ,$$

with $\eta \in \mathcal{H}$ arbitrary, continuous, and the **weak operator topology**, which is the coarsest topology for which all coefficient mappings

$$\mathcal{B}(\mathcal{H}) \ni T \mapsto \langle \varphi, T\eta \rangle \in \mathbb{C},$$

with $\varphi, \eta \in \mathcal{H}$ arbitrary, are continuous. Furthermore, let the **ultraweak topology** denote the coarsest topology for which all mappings

$$\mathcal{B}(\mathcal{H}) \ni T \mapsto \sum_{n\in\mathbb{N}} \langle \varphi_n, T\eta_n \rangle$$

are continuous. Here $(\eta_n)_{n\in\mathbb{N}}$ and $(\varphi_n)_{n\in\mathbb{N}}$ range over all families fulfilling

$$\sum_{n\in\mathbb{N}} \|\eta_n\|^2 < \infty \quad , \quad \sum_{n\in\mathbb{N}} \|\varphi_n\|^2 < \infty \quad .$$

We use the abbreviations **ONB** and **ONS** for orthonormal bases and orthonormal systems, respectively. $\dim(\mathcal{H})$ denotes the Hilbert space dimension, i.e., the cardinality of an arbitrary ONB of $\mathcal{H}$. Another abbreviation is the word **projection**, which in this book always refers to selfadjoint projection operators on a Hilbert space. For separable Hilbert spaces, the Hilbert space dimension is in $\mathbb{N} \cup \{\infty\}$, where the latter denotes the countably infinite cardinal. The standard index set of cardinality m (wherever needed) is $I_m = \{1, \ldots, m\}$, where $I_\infty = \mathbb{N}$, and the standard Hilbert space of dimension m is $\ell^2(I_m)$.

If $(\mathcal{H}_i)_{i\in I}$ is a family of Hilbert spaces, then $\bigoplus_{i\in I} \mathcal{H}_i$ is the space of vectors $(\varphi_i)_{i\in I}$ in the cartesian product fulfilling in addition

$$\|(\varphi_i)_{i\in I}\|^2 := \sum_{i\in I} \|\varphi_i\|^2 < \infty \quad .$$

The norm thus defined on $\bigoplus_{i\in I} \mathcal{H}_i$ is a Hilbert space norm, and $\bigoplus_{i\in I} \mathcal{H}_i$ is complete with respect to the norm. If the $\mathcal{H}_i$ are orthogonal subspaces of a common Hilbert space $\mathcal{H}$, $\bigoplus_{i\in I} \mathcal{H}_i$ is canonically identified with the closed subspace generated by the union of the $\mathcal{H}_i$.

If T is a densely defined operator on $\mathcal{H}$ which has a bounded extension, we denote the extension by $[T]$.

Unitary Representations

A **unitary, strongly continuous representation**, or simply **representation**, of a locally compact group G is a group homomorphism $\pi : G \to \mathcal{U}(\mathcal{H}_\pi)$ that is continuous, when the right hand side is endowed with the strong operator topology. Since weak and strong operator topology coincide on $\mathcal{U}(\mathcal{H}_\pi)$, the continuity requirement is equivalent to the condition that all coefficient functions of the type

$$G \ni x \mapsto \langle \varphi, \pi(x)\eta \rangle \in \mathbb{C},$$

are continuous.

Given representations σ, π, and operator $T : \mathcal{H}_\sigma \to \mathcal{H}_\pi$ is called **intertwining operator**, if $T\sigma(x) = \pi(x)T$ holds, for all $x \in G$. We write $\sigma \simeq \pi$ if σ and π are **unitarily equivalent**, which means that there is a *unitary* intertwining operator $U : \mathcal{H}_\sigma \to \mathcal{H}_\pi$. It is elementary to check that this defines an equivalence relation between representations. For any subset $\mathcal{K} \subset \mathcal{H}_\pi$ we let

$$\pi(G)\mathcal{K} = \{\pi(x)\eta : x \in G, \eta \in \mathcal{K}\} \quad .$$

A subspace of $\mathcal{K} \subset \mathcal{H}_\pi$ is called **invariant** if $\pi(G)\mathcal{K} \subset \mathcal{K}$. Orthogonal complements of invariant subspaces are invariant also. Restriction of a representation to invariant subspaces gives rise to **subrepresentations**. We write $\sigma < \pi$ if σ is unitarily equivalent to a subrepresentation of π. σ and π are called **disjoint** if there is no nonzero intertwining operator in either direction. A vector $\eta \in \mathcal{H}_\pi$ is called **cyclic** if $\pi(G)\eta$ spans a dense subspace of $\mathcal{H}_\pi$. A **cyclic representation** is a representation having a cyclic vector. All representations of interest to us are cyclic. In particular our standing assumption that G is second countable implies that all representations occurring in the book are realized on separable Hilbert spaces. π is called **irreducible** if every nonzero vector is cyclic, or equivalently, if the only closed invariant subspaces of $\mathcal{H}_\pi$ are $\{0\}$ and $\mathcal{H}_\pi$. Given a family $(\pi_i)_{i\in I}$, the direct sum $\pi = \bigoplus_{i\in I} \pi_i$ acts on $\bigoplus_{i\in I} \mathcal{H}_{\pi_i}$ via

$$\pi(x)\left(\varphi_i\right)_{i\in I} = \left(\pi_i(x)\varphi_i\right)_{i\in I} \quad .$$

The main result in connection with irreducible representations is Schur's lemma characterizing irreducibility in terms of intertwining operators. See [45, 3.5] for a proof.

Lemma 1.1. *If π_1, π_2 are irreducible representations, then the space of intertwining operators between π_1 and π_2 has dimension 1 or 0, depending on $\pi_1 \simeq \pi_2$ or not.*
In other words, π_1 and π_2 are either equivalent or disjoint.

Using the spectral theorem the following generalization can be shown. The proof can be found in [66, 1.2.15].

Lemma 1.2. *Let π_1, π_2 be representations of G, and let $T : \mathcal{H}_{\pi_1} \to \mathcal{H}_{\pi_2}$ be a closed intertwining operator, defined on a dense subspace $\mathcal{D} \subset \mathcal{H}_{\pi_1}$. Then $\overline{\mathrm{Im} T}$ and $(\ker T)^\perp$ are invariant subspaces and π_1, restricted to $(\ker T)^\perp$, is unitarily equivalent to the restriction of π_2 to $\overline{\mathrm{Im} T}$).*
If, moreover, π_1 is irreducible, T is a multiple of an isometry.

Given G, the **unitary dual** $\widehat{G}$ denotes the equivalence classes of irreducible representations of G. Whenever this is convenient, we assume the existence of a fixed choice of representatives of $\widehat{G}$, taking recourse to Schur's lemma to identify arbitrary irreducible representations with one of the representatives by means of the essentially unique intertwining operator.

We next describe the contragredient $\overline{\pi}$ of a representation π. For this purpose we define two involutions on $\mathcal{B}(\mathcal{H}_\pi)$, which are closely related to taking adjoints. For this purpose let $T \in \mathcal{B}(\mathcal{H}_\pi)$. If $(e_i)_{i \in I}$ is any orthonormal basis, we may define two linear operators T^t and $\overline{T}$ by prescribing

$$\langle T^t e_i, e_j \rangle = \langle T e_j, e_i \rangle \quad , \quad \langle \overline{T} e_i, e_j \rangle = \overline{\langle T e_i, e_j \rangle} \quad .$$

It is straightforward to check that these definitions do not depend on the choice of basis, and that $T^* = \overline{T}^t$, as we expect from finitedimensional matrix calculus. Additionally, the relations $\overline{T^t} = \overline{T}^t = T^*$ and $(ST)^t = T^t S^t$, $\overline{ST} = \overline{S}\,\overline{T}$ are easily verified.

Now, given a representation $(\pi, \mathcal{H}_\pi)$, the **(standard realization of the) contragredient representation** $\overline{\pi}$ acts on $\mathcal{H}_\pi$ by $\overline{\pi}(x) = \overline{\pi(x)}$. In general, $\overline{\pi} \not\simeq \pi$.

Commuting Algebras

The study of the commuting algebra, i.e., the bounded operators intertwining a representation with itself, is a central tool of representation theory. In this book, the **commutant** of a subset $M \subset \mathcal{B}(\mathcal{H})$, is denoted by M', and it is given by

$$M' = \{T \in \mathcal{B}(\mathcal{H}) : TS = ST \quad , \quad \forall S \in M\} \quad .$$

It is a **von Neumann algebra**, i.e. a subalgebra of $\mathcal{B}(\mathcal{H})$ which is closed under taking adjoints, contains the identity operator, and is closed with respect to the strong operator topology. The von Neumann density theorem [36, Theorem I.3.2, Corollary 1.3.1] states for selfadjoint subalgebras $\mathcal{A} \subset \mathcal{B}(\mathcal{H})$, that closedness in any of the above topologies on $\mathcal{B}(\mathcal{H})$ is equivalent to $\mathcal{A} = \mathcal{A}''$.

There are two von Neumann algebras associated to any representation π, the **commuting algebra** of π, which is the algebra $\pi(G)'$ of bounded operators intertwining π with itself, and the **bicommutant** $\pi(G)''$, which is the von Neumann algebra generated by $\pi(G)$. Since $\mathrm{span}(\pi(G))$ is a selfadjoint algebra, the von Neumann density theorem entails that it is dense in $\pi(G)''$ with respect to any of the above topologies. Invariant subspaces are conveniently discussed in terms of $\pi(G)'$, since a closed subspace $\mathcal{K}$ is invariant under π iff the projection onto $\mathcal{K}$ is contained in $\pi(G)'$.

Von Neumann algebras are closely related to the spectral theorem for selfadjoint operators, in the following way: Let $\mathcal{A}$ be a von Neumann algebra, and let T be a bounded selfadjoint operator. If S is an arbitrary bounded operator, it is well-known that S commutes with T iff S commutes with all spectral projections of T. Applying this to $S \in \mathcal{A}'$, the fact that $\mathcal{A} = \mathcal{A}''$ yields the following observation.

Theorem 1.3. *Let $\mathcal{A}$ is a von Neumann algebra on $\mathcal{H}$ and $T = T^* \in \mathcal{B}(\mathcal{H})$. Then $T \in \mathcal{A}$ iff all spectral projections of T are in $\mathcal{A}$.*

A useful consequence is that von Neumann algebras are closed under the functional calculus of selfadjoint operators, as described in [101, VII.7].

Corollary 1.4. *Let $\mathcal{A}$ is a von Neumann algebra on $\mathcal{H}$ and $T = T^* \in \mathcal{A}$ selfadjoint. Let $f : \mathbb{R} \to \mathbb{R}$ be a measurable function which is bounded on the spectrum of T. Then $f(T) \in \mathcal{A}$.*

Proof. Every spectral projection of $f(T)$ is a spectral projection of T. Hence the previous theorem yields the statement.

For more details concerning the spectral theorem we refer the reader to [101, Chapter VII]. The relevance of the spectral theorem for the representation theory of the reals is sketched in Section 2.7.

Tensor Products

The tensor product notation is particularly suited to treating direct sums of equivalent representations. Let $\mathcal{H}, \mathcal{K}$ be Hilbert spaces. The **Hilbert space tensor product** $\mathcal{H} \otimes \mathcal{K}$ is defined as the space of bounded linear operators $T : \mathcal{K} \to \mathcal{H}$ satisfying

$$\|T\|^2_{\mathcal{H}\otimes\mathcal{K}} := \sum_{j\in J} \|Te_j\|^2 < \infty \quad .$$

Here $(e_j)_{j\in J}$ is an ONB of $\mathcal{K}$. The Parseval equality can be employed to show that the norm is independent of the choice of basis, making $\mathcal{H} \otimes \mathcal{K}$ a Hilbert space with scalar product

$$\langle S, T\rangle = \sum_{j\in J} \langle Se_j, Te_j\rangle \quad .$$

Of particular interest are the operators of rank one. We define the **elementary tensor** $\varphi \otimes \eta$ as the rank one operator $\mathcal{K} \to \mathcal{H}$ defined by $\mathcal{K} \ni z \mapsto \langle z, \eta\rangle\varphi$. The scalar product of two rank one operators can be computed as

$$\langle \eta \otimes \varphi, \eta' \otimes \varphi'\rangle_{\mathcal{H}\otimes\mathcal{K}} = \langle \eta, \eta'\rangle_{\mathcal{H}} \langle \varphi', \varphi\rangle_{\mathcal{K}} \quad .$$

Note that our definition differs from the one in [45] in that our tensor product consists of *linear* operators as opposed to *conjugate-linear* in [45]. As a consequence, our elementary tensors are only conjugate-linear in the $\mathcal{K}$ variable, as witnessed by the change of order in the scalar product. However, the arguments in [45] are easily adapted to our notation. For computations in $\mathcal{H} \otimes \mathcal{K}$, it is useful to observe that ONB's $(\eta_i)_{i \in I} \subset \mathcal{H}$ and $(\varphi_j)_{j \in J} \subset \mathcal{K}$ yield an ONB $(\eta_i \otimes \varphi_j)_{i \in I, j \in J}$ of $\mathcal{H} \otimes \mathcal{K}$ [45, 7.14]. By collecting terms in the expansion with respect to the ONB, one obtains that each $T \in \mathcal{H} \otimes \mathcal{K}$ can be written as

$$T = \sum_{j \in J} a_j \otimes \varphi_j = \sum_{i \in I} \eta_i \otimes b_i \quad , \tag{1.5}$$

where the a_j and b_i are computed by $a_j = T\varphi_j$ and $b_i = T^* \eta_i$, yielding

$$T = \sum_{j \in J} (T\varphi_j) \otimes \varphi_j = \sum_{i \in I} \eta_i \otimes (T^* \eta_i) \quad , \tag{1.6}$$

as well as

$$\|T\|_2^2 = \sum_{j \in J} \|a_j\|^2 = \sum_{i \in I} \|b_i\|^2 \quad .$$

By polarization of this equation we find that given a second operator $S = \sum_{j \in J} c_j \otimes \varphi_j$, the scalar product can also be computed via

$$\langle T, S \rangle = \sum_{j \in J} \langle a_j, c_j \rangle \quad .$$

Operators $T \in \mathcal{B}(\mathcal{H})$, $S \in \mathcal{B}(\mathcal{K})$ act on elements on $\mathcal{H} \otimes \mathcal{K}$ by multiplication. On elementary tensors, this action reads as

$$(T \otimes S)(\eta \otimes \varphi) = T \circ (\eta \otimes \varphi) \circ S = (T\eta) \otimes (S^* \varphi) \quad ,$$

which will be denoted by $T \otimes S \in \mathcal{B}(\mathcal{H} \otimes \mathcal{K})$. Keep in mind that this tensor is also only sesquilinear. Given two representations π, σ, the **tensor product representation** $\pi \otimes \overline{\sigma}$ is the representation of the direct product $G \times G$ acting on $\mathcal{H}_\pi \otimes \mathcal{H}_\sigma$ via $\pi \otimes \overline{\sigma}(x, y) = \pi(x) \otimes \sigma(y)^*$. On elementary tensors this action is given by

$$(\pi \otimes \overline{\sigma}(x, y))(\eta \otimes \varphi) = (\pi(x)\eta) \otimes (\sigma(x)\varphi) \quad .$$

Observe that the sesquilinearity of our tensor product notation entails that the restriction of $\pi \otimes \overline{\sigma}$ to $\{1\} \otimes G$ is indeed equivalent to $\dim(\mathcal{H}_\pi) \cdot \overline{\sigma}$, where $\overline{\sigma}$ is the contragredient of σ.

One can use the tensor product notation to define a compact realization of the multiple of a fixed representation. Given such a representation σ, the standard realization of $\pi = m \cdot \sigma$ acts on $\mathcal{H}_\pi = \mathcal{H}_\sigma \otimes \ell^2(I_m)$ by

$$\pi(x) = \sigma(x) \otimes \mathrm{Id}_{\ell^2(I_m)} \quad .$$

The advantage of this realization lies in compact formulae for the associated von Neumann algebras, if σ is irreducible:

$$\pi(G)' = 1 \otimes \mathcal{B}(\ell^2(I_m)) \quad , \tag{1.7}$$

, which is understood as the algebra of all operators of the form $\mathrm{Id}_{\mathcal{H}_\sigma} \otimes T$, and

$$\pi(G)'' = \mathcal{B}(\mathcal{H}_\sigma) \otimes 1 \quad , \tag{1.8}$$

with analogous definitions. The follow for instance by [105, Theorem 2.8.1].

Trace Class and Hilbert-Schmidt Operators

Given a bounded positive operator T on a separable Hilbert space $\mathcal{H}$, T it is called **trace class operator** if its **trace class norm**

$$\|T\|_1 = \mathrm{trace}(T) = \sum_{i \in I} \langle T\eta_i, \eta_i \rangle < \infty \quad ,$$

where $(\eta_i)_{i \in I}$ is an ONB of $\mathcal{H}$. $\|T\|_1$ can be shown to be independent of the choice of ONB. An arbitrary bounded operator T is a **trace class operator** iff $|T|$ is of trace class. This defines the Banach space $\mathcal{B}_1(\mathcal{H})$ of trace clase operators. The trace

$$\mathrm{trace}(T) = \sum_{n \in \mathbb{N}} \langle T\eta_i, \eta_i \rangle$$

is a linear functional on $\mathcal{B}_1(\mathcal{H})$, and again independent of the choice of ONB. A useful property of the trace is that $\mathrm{trace}(TS) = \mathrm{trace}(ST)$, for all $T \in \mathcal{B}_1(\mathcal{H})$ and $S \in \mathcal{B}_\infty(\mathcal{H})$.

More generally, we may define for arbitrary $1 \le p < \infty$ the **Schatten-von Neumann space of order** p as the space $\mathcal{B}_p(\mathcal{H})$ of operators T such that $|T|^p$ is trace class, endowed with the norm

$$\|T\|_p = \|(T^*T)^{p/2}\|_1^{1/p} \quad .$$

Again $\mathcal{B}_p(\mathcal{H})$ is a Banach space with respect to $\| \cdot \|_p$. An operator T is in $\mathcal{B}_p(\mathcal{H})$ iff $|T|$ has a discrete p-summable spectrum (counting multiplicities). This also entails that $\mathcal{B}_p(\mathcal{H}) \subset \mathcal{B}_r(\mathcal{H})$, for $p \le r$, and that these spaces are contained in the space of compact operators on $\mathcal{H}$. Moreover, it entails that $\| \cdot \|_\infty \le \| \cdot \|_p$.

As a further interesting property, $\mathcal{B}_p(\mathcal{H})$ is a twosided ideal in $\mathcal{B}(\mathcal{H})$, satisfying

$$\|ATB\|_p \le \|A\|_\infty \|T\|_p \|B\|_\infty \quad .$$

We will exclusively be concerned with $p = 1$ and $p = 2$. Elements of the latter space are called **Hilbert-Schmidt operators**) . $\mathcal{B}_2(\mathcal{H})$ is a Hilbert space, with scalar product

$$\langle S, T\rangle = \mathrm{trace}(ST^*) = \mathrm{trace}(T^*S) \quad .$$

In fact, as the formula

$$\mathrm{trace}(T^*T) = \sum_{i\in I} \|T\eta_i\|^2 = \|T\|^2_{\mathcal{H}\otimes\mathcal{H}}$$

shows, $\mathcal{B}_2(\mathcal{H}) = \mathcal{H}\otimes\mathcal{H}$. In particular, all facts involving the role of rank-one operators and elementary tensors presented in the previous section hold for $\mathcal{B}_2(\mathcal{H})$.

Measure Spaces

In this book integration, either on a locally compact group or its dual, is ubiquitous. Borel spaces provide the natural context for our purposes, and we give a sketch of the basic notions and results. For a more detailed exposition, confer the chapters dedicated to the subject in [15, 17, 94].

Let us quickly recall some definitions connected to measure spaces. A **Borel space** is a set X equipped with a σ-algebra $\mathcal{B}$, i.e. a set of subsets of X (containing the set X itself) which is closed under taking complements and countable unions. $\mathcal{B}$ is also called **Borel structure**. Elements of $\mathcal{B}$ are called **measurable** or **Borel**. A σ-algebra **separates points**, if it contains the singletons. Arbitrary subsets A of a Borel space $(X, \mathcal{B})$, measurable or not, inherit a Borel structure by declaring the intersections $A\cap B$, $B\in\mathcal{B}$, as the measurable sets in A.

In most cases we will not explicitly mention the σ-algebra, since it is usually provided by the context. For a locally compact group, it is generated by the open sets. For countable sets, the power set will be the usual Borel structure. A **measure space** is a Borel space with a (σ-additive) measure μ on the σ-algebra. **ν-nullsets** are sets A with $\nu(A) = 0$, whereas **conull sets** are complements of nullsets.

If μ and ν are measures on the same space, μ is **ν-absolutely continuous** if every ν-nullset is a μ-nullset as well. We assume all measures to be σ-finite. In particular, the Radon-Nikodym Theorem holds [104, 6.10]. Hence absolute continuity of measures is expressable in terms of densities.

Measurable mappings between Borel spaces are defined by the property that the preimages of measurable sets are measurable. A bijective mapping $\phi : X \to Y$ between Borel spaces is a **Borel isomorphism** iff ϕ and ϕ^{-1} is measurable. A mapping $X \to Y$ is **μ-measurable** iff it is measurable outside a μ-nullset. For complex-valued functions f given on any measure space, we let $\mathrm{supp}(f) = f^{-1}(\mathbb{C}\setminus 0)$. Inclusion properties between supports are understood to hold only up to sets of measure zero, which is reasonable if one deals with L^p-functions. Given a Borel set A, we let $\mathbf{1}_A$ denote its indicator function.

Given a measurable mapping $\varPhi : X \to Y$ between Borel spaces and a measure μ on X, the image measure $\varPhi^*(\mu)$ on Y is defined as $\varPhi^*(\mu)(A) =$

$\mu(\Phi^{-1}(A))$. A measure ν on Y is a **pseudo-image** of μ under Φ if ν is equivalent to $\Phi^*(\tilde{\mu})$, and $\tilde{\mu}$ is a finite measure on X which is equivalent to μ. $\tilde{\mu}$ exists if μ is σ-finite. Clearly two pseudo-images of the same measure are equivalent.

Let us now turn to locally compact groups G and G-spaces. A **G-space** is a set X with a **an action of** G **on** X, i.e., a mapping $G \times X \to X$, $(g, x) \mapsto g.x$, fulfilling $e.x = x$ and $g.(h.x) = (gh).x$. A **Borel G-space** is a G-space with the additional property that G and X carry Borel structures which make the action measurable; here $G \times X$ is endowed with the product Borel structure. If X is a G-space, the **orbits** $G.x = \{g.x : g \in G\}$ yield a partition of X, and the set of orbits or **orbit space** is denoted X/G for the orbit space. This notation is also applied to invariant subsets: If $A \subset X$ is **G-invariant**, i.e. $G.A = A$, then A/G is the space of orbits in A, canonically embedded in X/G. If X is a Borel space, the **quotient Borel structure** on X is defined by declaring all subsets $A \subset X/G$ as Borel for which the corresponding invariant subset of X is Borel. It is the coarsest Borel structure for which the quotient map $X \to X/G$ is measurable.

For $x \in X$ the **stabilizer of** x is given by $G_x = \{g \in G : g.x = x\}$. The canonical map $G \ni g \mapsto g.x$ induces a bijection $G/G_x \to G.x$.