Underwater object classification using scattering transform of sonar signals

Naoki Saito^a and David S. Weber^a

^aUC Davis, One Shields Ave, Davis, CA, USA

ABSTRACT

In this paper, we apply the scattering transform (ST)—a nonlinear map based off of a convolutional neural network (CNN)—to classification of underwater objects using sonar signals. The ST formalizes the observation that the filters learned by a CNN have wavelet-like structure. We achieve effective binary classification both on a real dataset of Unexploded Ordinance (UXOs), as well as synthetically generated examples. We also explore the effects on the waveforms with respect to changes in the object domain (e.g., translation, rotation, and acoustic impedance, etc.), and examine the consequences coming from theoretical results for the scattering transform. We show that the scattering transform is capable of excellent classification on both the synthetic and real problems, thanks to having more quasi-invariance properties that are well-suited to translation and rotation of the object.

Keywords: Scattering Transform, Synthetic Aperature Sonar, LASSO, Underwater Object Classification

1. INTRODUCTION

Deep neural networks, and convolutional neural networks in particular, have proven quite effective at discerning hierarchical patterns in large datasets.¹ Some examples include image classification,² face recognition,³ and speech recognition,⁴ among many others. Clearly, something is going very right in the design of CNNs. However, the principles that account for this success are still somewhat elusive,⁵ as is the construction of systems that work well with few examples.

The scattering transform was created to remedy these issues. In 2012, Stephane Mallat and Joan Bruna published both theoretical results⁶ and numerical implementations⁷ tying together convolutional neural networks and wavelet theory. They demonstrated that scattering networks of wavelets and modulus nonlinearities, are translation invariant in the limit of infinite scale, and Lipschitz continuous under non-uniform translation, i.e. $T_{\tau}(f)(x) = f(x - \tau(x))$ for τ with bounded gradient. Numerically, they achieved state of the art on image and texture classification problems.

More recent work from Wiatowski and Bölcskei have generalized the Lipschitz continuity result from wavelet transforms to frames, and more importantly, established that increasing the *depth* of the network also leads to translation invariant features.⁸ There have been a number of follow up papers, including a discrete version of Wiatowski's result,⁹ a related method on graphs,¹⁰ and a pseudo-inverse using phase retrieval.¹¹ There have also been a number of papers using the scattering transform in such problems as fetal heart rate classification,¹² age estimation from face images,¹³ and voice detection in the presence of transient noise.¹⁴

In this paper, we will apply the scattering transform of Mallat et al. to the problem of classification of underwater objects using sonar. This is motivated by previous work done by our group on the detection of unexploded ordinance (UXO) using synthetic aperture sonar (SAS).¹⁶ From this, we use the BAYEX14 dataset,¹⁵ which contains 14 different objects at various distances and rotations, partially submerged in a thin layer of mud on a sandy shallow ocean bed off Florida. Fig. 1 shows examples of these objects. Additionally, we can generate synthetic examples using a fast solver for the Helmholtz equation in two regions with differing speed of sound,

Further author information: (Send correspondence to David S. Weber)

Naoki Saito: E-mail: saito@math.ucdavis.edu,

website: https://www.math.ucdavis.edu/~saito/

David S. Weber.: E-mail: dsweber@math.ucdavis.edu,

website: https://dsweber2.wordpress.com/



Figure 1: Example targets on shore and in the target environment¹⁵

provided by Ian Sammis and James Bremer.^{17,18} We use this to examine more closely the dependence of both classification and the output of the scattering transform on both material properties and shape variations.

The approximate translation invariance and Lipschitz continuity under small diffeomorphisms of the scattering transform mean that for classes that are invariant under these transformations, members of the same class will be close together in the resulting space. So long as morphing via $f(t - \tau(t))$ from one class to another requires a τ with large derivative, then the classes will be well separated. Accordingly, we use a linear classifier on the output of the scattering transform. Additionally, because the scattering transform concentrates energy at coarser scales and wavelets in general encourage sparsity for smooth signals with singularities, we use a sparse version of logistic regression as our linear classifier, using LASSO.¹⁹ Specifically, we use a julia wrapper around GLMNET.²⁰

As a baseline classification scheme to compare the classification performance, we use the same linear classifier on the absolute value of the Fourier transform (AVFT) of the signal. One strong advantage of the absolute value of the Fourier transform is completely invariant to translations, and is the prototype of every linear filter which is invariant to translations.^{21, 22} In addition to this invariant, the close ties between frequency and the speed of sound suggest that it should be sensitive to changes in the material. This will be examined in more depth in Section 4.

Section 2 gives the basic set-up of a scattering transform, and approximate invariants from previous theoretical developments. These are increasing translation invariance with increasing depth, and Lipschitz stability to nonuniform time and frequency translation. Section 3 describes the setup used to generate both the synthetic and real datasets. Section 5 gives the results of applying the ST and AVFT to the dataset to performing binary classification on shape and material in the synthetic case, and between UXO's and various other objects in the real case.

2. GENERALIZED SCATTERING TRANSFORM

A generalized scattering network (hereafter referred to simply as a scattering network or ST),^{6,8} has an architecture that is a continuous analog of a CNN. For a diagram, see Fig. 2. First, at layer m, we start with a family of generators $\{g_{\lambda(m)}\}_{\lambda(m)\in\Lambda_m} \subset \mathcal{L}^1(\mathbb{R}^d) \cap \mathcal{L}^2(\mathbb{R}^d)$ for a translation invariant frame $\Psi_m = \{\psi_{b,\lambda(m)}\}_{b\in\mathbb{Z}^d,\lambda(m)\in\Lambda_m} =$

 ${T_b Ig_{\lambda^{(m)}}}_{b \in \mathbb{Z}^d, \lambda^{(m)} \in \Lambda_m}$ with frame bounds a_m and b_m , that is

$$a_m \|f\|_2^2 \leq \sum_{\lambda^{(m)} \in \Lambda_m, b \in \mathbb{Z}^d} \left| \langle T_b C g_{\lambda^{(m)}}, f \rangle \right|^2 \leq b_m \|f\|_2^2.$$

for all $f \in \mathcal{L}^2(\mathbb{R}^d)$. Here, $T_t[f](x) = f(x-t)$ is the translation operator, $C[f](x) = \overline{f(-x)}$ is the involution operator, and Λ_m is some countable discrete index set, such as \mathbb{Z} . This index set needs to tile the frequency plane in some way, for example by indexing scales and rotations. The frame atoms $g_{\lambda^{(m)}}$ correspond to the receptive fields found in each layer of a CNN. In addition, at layer m we define a Lipschitz-continuous operator M_m with bound γ_m which satisfies $M_m f = 0 \Rightarrow f \equiv 0$. After these have been applied, the result is sub-sampled at a rate $r_m \geq 1$. Finally, the output at each layer is then generated by averaging with a specific atom $\chi_m = g_{\lambda_*^{(m)}}$, which is removed from the set of frame atoms.

The scattering transform of Mallat⁶ for \mathbb{R}^d corresponds to choosing $\Lambda_m = \{2^{j/Q_m}h\}_{j>-J_m,h\in H_m}$ for some finite rotation group H_m , $g_{0,1} = \psi(x)$ for some mother wavelet ψ , $\chi_m = \phi$ its corresponding father wavelet, and Q_m a quality factor. The Lipschitz nonlinearity is $|\cdot|$. However, it lacks a sub-sampling factor, so $r_m = 1$. Explicitly, the generator corresponding to index $\lambda^{(m)} = (j,h)$ is $g_{\lambda^{(m)}} = 2^{dj/Q_m} \psi(2^{j/Q_m} h^{-1}x)$.

To get from layer m-1 to layer m, we define the function $u_m : \Lambda_m \times \mathcal{L}^2(\mathbb{R}^d) \to \mathcal{L}^2(\mathbb{R}^d)$ by

$$u_m[\lambda^{(m)}](f)(z) = M_m[f \star g_{\lambda^{(m)}}](r_m z).$$

Using this to define the value at a layer m, we have a path coming along a path of indices $q \in \Lambda^m := \Lambda_m \times \ldots \times \Lambda_1$,

$$u[q](f) = u_m[\lambda^{(m)}]u_{m-1}[\lambda^{(m-1)}]\cdots u_1[\lambda^{(1)}]f.$$
(1)

Choosing d = 1 corresponds to audio signals, such as sonar. Unlike a CNN, in a Scattering transform each layer has output, including just the average $\chi_m \star f$, which is layer zero. For m > 1, the output is

$$\Phi_m[f] \coloneqq \left\{ u[q]f \star \chi_m \right\}_{q \in \Lambda^m}$$

while for the entire scattering transform,

$$\Phi[f] \coloneqq \bigcup_{m=0}^{\infty} \Phi_m[f].$$

2.1 Previous theory

There are two additional conditions that restrict the various operators in a given layer simultaneously.⁸ The first is the *weak admissibility condition*, which requires that the upper frame bound b_m be sufficiently small compared to the sub-sampling factor and Lipschitz constants:

$$\max\left\{b_m, \frac{b_m \gamma_m^2}{r_m^d}\right\} \le 1,$$

which can be achieved by scaling the $g_{\lambda^{(m)}}$.

The second is that the nonlinearities must commute with the translation operator, so $M_m T_t[f] = T_t M_m[f]$. Most nonlinearities used for CNN's are pointwise i.e., $M_m[f](x) = \rho_m(f(x))$, so they certainly commute with T_t . Given these constraints, we can now state the results of Wiatowski and Bölcskei⁸ precisely; the first is that the resulting features Φ_m deform stably with respect to small frequency and space deformations:

THEOREM 1.⁸ For frequency shift $\omega \in C(\mathbb{R}^d; \mathbb{R})$ and space shift $\tau \in C^1(\mathbb{R}^d)$, define the operator $F_{\tau,\omega}[f](x) \coloneqq e^{2\pi i \omega(x)} f(x - \tau(x))$. If $\|D\tau\|_{\infty} \leq \frac{1}{2d}$, then there exists a C > 0 independent of the choice of parameters for Φ s.t. for all $f \in \mathcal{L}^2_a(\mathbb{R}^d)$,

$$\left\|\Phi[F_{\tau,\omega}f] - \Phi[f]\right\|_{2} \coloneqq \sum_{m=1}^{\infty} \sum_{q \in \Lambda^{m}} \left\|\chi_{m} \star u[q](F_{\tau,\omega}f) - \chi_{m} \star u[q]f\right\|_{2} \le C(r\|\tau\|_{\infty} + \|\omega\|_{\infty})\|f\|_{2}, \tag{2}$$



Figure 2: Generalized scattering transform. Here, $q_{k_n,\ldots,k_0} = (\lambda_{k_n}^{(n)},\ldots,\lambda_{k_1}^{(1)})$ is an element of Λ^n , so $u_n[q_{k_m,\ldots,k_0}]$ is as in Eq. (1), and $s_n[q_{k_n,\ldots,k_0}] = \chi_n \star u_n[q_{k_m,\ldots,k_0}]f$, i.e. an element of Φ_n

where $\mathcal{L}^2_a(\mathbb{R}^d)$ is the set of $\mathcal{L}^2(\mathbb{R}^d)$ functions whose Fourier transforms are band limited to [-a, a]. Mallat shows a similar bound for the specific case that $M_m = |\cdot|$ and $g_{\lambda(m)}$ is generated by an admissible mother wavelet ψ with a number of conditions.⁶

Their next result, translation invariance that increases with depth, is distinct from the one shown by Mallat where translation invariance increases with resolution:

THEOREM 2.⁸ Given the conditions above, for $f \in \mathcal{L}^2(\mathbb{R}^d)$ the mth layer's features satisfy

$$\Phi_m[T_t f] = T_{\frac{t}{r_1 \cdots r_m}} [\Phi_m(f)].$$

If there is also a global bound K on the decay of the Fourier transforms of the output features χ_n :

$$|\widehat{\chi_m}(\omega)| \le K_{\varepsilon}$$

then we have the stronger result

$$\sum_{n=1}^{m} \left\| \Phi_n[T_t f] - \Phi_n[f] \right\|_2 \le \frac{2\pi |t| K}{r_1 \cdots r_m} \|f\|_2.$$

To compare with the result from Mallat,⁶ that says that for an admissible mother wavelet ψ , the scattering transform achieves perfect translation invariance as the lower bound on the scale goes to infinity:

$$\lim_{J \to \infty} \left\| \Phi_J[T_t f] - \Phi_J[f] \right\| = 0.$$

3. SONAR DETECTION

The problem that we will be investigating with the scattering transform techniques is the classification of objects partially buried on the sea floor using sonar data. This is motivated by using an unmanned underwater vehicle (UUV) equipped with sonar to detect unexploded ordinance. For this problem, we have both real and synthetic examples. The real examples consist of 14 partially buried objects at various distances and rotations, in a shallow mud layer on top of a sand ocean bed at a depth of $\sim 8m$. BAYEX14 simulated the path of a UUV by stationing a sensor/emitter onto a rail, and pinging the field of objects was pinged at short intervals, see Fig. 3. So for each rotation of the object relative to the rail, there is a 2D wavefield, where each signal corresponds to a location on the rail and the observation time.



Figure 3: The three shapes of Ω used in the synthetic setting. The rectangle on the left includes the observation rail without rotation. The range is approximately 10m, while the observation rail itself is 12m in total. The triangle has side lengths of 1m, while the shark-fin is deformed from the same triangle. The rectangle has side lengths of 1m and 1.5m.

The synthetic examples come from considering the 2D Helmholtz equation in the regions with differing speed of sound:¹⁷

$$\begin{aligned} \Delta u + k_1^2 u &= 0 \quad \text{in } \Omega \\ \Delta v + k_2^2 v &= 0 \quad \text{in } \Omega^c \\ u - v &= g \quad \text{on } \partial\Omega \\ \partial_\nu u - \partial_\nu v &= \partial_\nu g \quad \text{on } \partial\Omega \\ \sqrt{|x|} \left(\partial_{|x|} - \mathrm{i}k_2 \right) v(x) \to 0 \text{ as } |x| \to \infty, \end{aligned}$$

which gives the response to a sinusoidal signal with frequency ω on an object with $k_i = \omega/c_i$, where c_i is the speed of sound in the material, ranging from 343m/s in air, to 1503m/s in water, to 5100m/s in aluminum. It is worth noting that this is idealized in several ways: the model we use is 2D, rather than 3D, the material is modeled as a fluid with only one layer, instead of a solid with multiple different components, and there is no representation of the ocean floor itself.

The signals sent out by the UUVs are not pure sinusoids. One can approximate the response to multi-frequency signals (e.g. Gabor functions or chirps) by integrating across frequencies. We use a fast solver created by Ian Sammis and James Bremer to synthesize a set of examples, where we can more explicitly test the dependence of the scattering transform on the material properties (corresponding to changing the speed of sound) and geometry. The current dataset we use was created by Vincent Bodin, a former summer intern supervised by the first author. The input/source signal is shown in Fig. 5; zero padding is applied to each source signal to make the periodized version have approximately the same behavior as an input with finite support.

There are many design choices in testing the effectiveness of a classifier. As a baseline to compare against, we use the same **GLMNET** logistic classifier on the absolute value of the Fourier transform (AVFT), which is a simple classification technique that is translation invariant and sensitive to frequency shifts. To understand the generalization ability of the techniques, we split the data into two halves, one half training set and one half test set, uniformly at random 10 times, i.e. 10-fold cross validation. For the synthetic dataset, we also normalized the signals and added uniform Gaussian white noise so the SNR is 5dB.

For the synthetic data, there are two primary problems of interest. The first is determining the effects of varying shape on the scattering transform. An example of the 1D scattering transform for a triangle is in Fig. 4. Since the energy at each layer decays exponentially with layer index, layers 1-3 have been scaled to match the intensity of the first layer. Note that only the zeroth layer has negative values; this is because the nonlinearity used by the scattering transform is an absolute value. In the figure, one can clearly see a time concentrated portion of the signal in layers 0, 1, and 2.

For the real dataset, the problem of interest is somewhat more ambiguous. In addition to a set of UXO's and a set of arbitrary objects, there are some UXO replicas, not all of which are made of the same material. As we will see in the synthetic case, the difference in the speed of sound has a much clearer effect on classification



Figure 4: The ST coefficients for the triangle object filled with a material whose speed of sound is $c_1 = 2000$ m/s, while the ambient speed of sound is $c_2 = 1503$ m/s. The x-axis refers to the depth, varying from m = 0 to m = 3. Scale is coarsest at the highest y-values, while within a given layer, the time increases horizontally. Note that only paths of increasing scale have been kept for computational reasons.

accuracy than shape, so it is somewhat ambiguous how to treat these. We should expect that correctly classifying non-UXO's to be more difficult, since as a class they do not have much in common– a SCUBA tank bears more resemblance to a UXO than it does to a rock.



4. GEOMETRIC PROPERTIES

Ideally the invariants discussed above would apply to transformations in the *object* domain rather than to signals. But translations of the object (or equivalently, the observation rail), have a more complicated effect on the signal than simply translating the observation; at a minimum translating away from the object will cause a decay in signal amplitude. The changes in the object domain we seek to understand are changes in object material, translations and rotations of the object/rail, and changes in geometry. A classifier for this problem should be invariant to translation and rotation, but sensitive to the geometry of the object and the material.

To understand more deeply how changes in the geometry affect the observed signals, we need to examine the solutions more closely. Let the input signal be fixed as $s(t) \in \mathcal{L}^1(\mathbb{R}) \cap \mathcal{L}^2(\mathbb{R})$. The ideal reconstruction of the response to s, for a transmitter and receiver located at $\vec{x} \in \Omega^c$ is

$$f(t, \vec{x}) = \int_{-\infty}^{\infty} \widehat{s}(\omega) v(\omega, \vec{x}) e^{-i\omega \left(t - P(\omega, \vec{x})\right)} d\omega$$
(3)

where v is the solution, \hat{s} denotes the Fourier transform of s. To define the observation rail, first we define an unrotated observation, $\vec{x}_0(r) = (x, r)$ for $r \in [-y, y]$, so the rail has range x = 10m and length 2y = 12m. Then, a rotation of the object by an angle $-\theta$ is equivalent to rotating the rail by θ , so define $\vec{x}_{\theta}(r) = R_{\theta}\vec{x}_0(r)$ where R_{θ} is the appropriate rotation matrix. In this setup, x = 10m, and y = 6m.



Figure 6: Non-zero portion of the rectangle signal for varying speed of sound, when facing the long edge.

4.1 Effects of the speed of sound

To determine the behavior of a fixed location $x_{\theta}(r)$ as we change the speed of sound, we use common acoustic properties such as reflection coefficients and Snell's law. The crudest possible assumption that still gives meaningful results is that internal angles are irrelevant, and only refraction, reflection, and distance matters. Going from Ω^c with speed of sound c_2 to Ω with speed of sound c_1 , the reflection coefficient is given by $V_{2,1} = \frac{Z_2 - Z_1}{Z_1 + Z_2}$, while the refraction coefficient is $W_{2,1} = 1 - V_{2,1} = \frac{2Z_1}{Z_1 + Z_2}$, where the impedance is $Z_i = \rho_i c_i$ (ρ_i is the density of the material). The distance from the center to a given point $x_{\theta}(r)$ on the line is just given by the Pythagorean theorem, $\sqrt{x^2 + r^2}$. This means the initial peak occurs at $\frac{1}{c_2}\sqrt{x^2 + r^2}$. If the input peak has magnitude A_0 , then the first return peak should be approximately $A_1 = V_{2,1}\frac{A_0}{x^2 + r^2} = \frac{(Z_2 - Z_1)}{(x^2 + r^2)(Z_2 + Z_1)}A_0$; since $Z_2 > Z_1$ for most relevant examples, this is positive. Setting diam(Ω) = D, the next peak can be approximated as $A_2 = W_{2,1}V_{1,2}W_{1,2}\frac{A_0}{(x^2 + r^2)D^2} = \frac{4Z_2Z_1(Z_1 - Z_2)}{(x^2 + r^2)D(Z_1 + Z_2)^3}A_0$; the sign of the second peak will flip. The third peak is $A_3 = W_{2,1}V_{1,2}^2W_{1,2}\frac{A_0}{(x^2 + r^2)D^2}$ and will flip sign again. Similarly $A_n = W_{2,1}V_{1,2}^nW_{1,2}\frac{A_0}{(x^2 + r^2)D^n}$.

From this, we should expect that the decay from the first peak to the second peak, $\frac{4Z_1Z_2}{D(Z_1+Z_2)^2}$, is larger than that between any further consecutive peaks, $\frac{2|Z_1-Z_2|}{(Z_1+Z_2)D}$. One can see the sign flip clearly in Fig. 6 (take care that the input signal (Fig. 5) is a positive spike followed by a negative spike, so the second peak begins at ~ 75s for speed of sound c_1 m/s).

For classification purposes, for a fixed angle and position on the rail, the reflection of every peak but the first depends on both D and the speed of sound c_1 . Further, the time between peaks should be D/c_1 , so the scale of the solutions should provide a strong indicator of the speed of sound in the material. We will indeed see that the absolute value of a Fourier transform (AVFT), which has access to scale information, is reasonably effective at this problem. However for shapes with approximately the same diameter, such as the triangle and shark-fin, none of the above is useful for discrimination.

4.2 Effects of rotation

The discussion so far has almost completely ignored the internal geometry of Ω . For a rectangle, when facing the longer side, this approximation is reasonable, as Snell's law matters when the internal angles depart significantly from $\pi/2$. However, the problem of ray tracing in a region Ω is non-trivial in its own right, requiring averaging over all paths.

To avoid this issue, instead of trying to directly construct properties of the observations, we can derive how they will change under rotation and translation. We can use the far field approximation to do so [23, Chpt. 4]. Since diam(Ω) \approx 1 and the center frequency of s is $\omega_0 = 2500$ Hz, we have that the range $10m \gg \frac{c_1}{\omega_0} \approx .6m$, the condition for far-field. Here we take the solution as approximately separable, so $v(\omega, \vec{x})e^{i\omega P(\omega, \vec{x})} \approx R(\omega, r)\Theta(\omega, \theta)$. In this representation, R is a Bessel function of the first kind, and so to zero-th order is approximately $R(\omega, r) = J_0(k_2r) \approx \frac{1}{\sqrt{k_2r}}e^{ik_2r-\pi/4}$.



Figure 7: Composing a rail observation along $\vec{x}_0(q)$ from that of $\vec{x}_{-1}(p)$ and $\vec{x}_1(r)$.

Using this, we can examine the effect of rotation of the object (or the rail about the object). Suppose we know the solution at angles θ_{-1} and θ_1 , and we want to determine the solution at an angle θ_0 between these. Every point in $\vec{x}_0(q)$ will have the same angle as a point on either $\vec{x}_1(r)$ or $\vec{x}_{-1}(p)$ (or possibly both) if θ_{-1} and θ_1 are close enough that the paths cross, such as the case in Fig. 7. This happens when $\tan(\frac{1}{2}(\theta_1 - \theta_{-1})) < y/x$, or in the synthetic dataset, $\theta_1 - \theta_{-1} \le \pi/6$). Some geometry gives that the point $\vec{x}_0(q)$ has the same angle as $\vec{x}_1(r)$ if

$$T(q) = -x \tan(\theta_1 - \theta_0 - \arctan(q/x))$$
(4)

for q > 0; similar reasoning works for q < 0 and \vec{x}_{-1} with flipped signs. The distance changes from $\sqrt{x^2 + q^2}$ to $\sqrt{x^2 + T(q)^2}$, while the angle remains fixed, so plugging the zero-th order approximation to the Bessel function

 $J_0(k_2r)$ above into Eq. (3),

$$f(r, \vec{x}_{0}(q)) \approx \int_{-\infty}^{\infty} \widehat{s}(\omega) \frac{J_{0}\left(k_{2}\sqrt{q^{2}+x^{2}}\right)}{J_{0}\left(k_{2}\sqrt{T(q)^{2}+x^{2}}\right)} v\left(\omega, \vec{x}_{1}\left(T(q)\right)\right) e^{-i\omega(t-P(\omega, \vec{x}_{1}(T(q))))} d\omega$$
(5)
$$f(t, \vec{x}_{0}(q)) \approx \frac{\sqrt[4]{T(q)^{2}+x^{2}}}{\sqrt[4]{q^{2}+x^{2}}} \int_{-\infty}^{\infty} \widehat{s}(\omega) v\left(\omega, \vec{x}_{1}\left(T(q)\right)\right) e^{-i\omega\left(t+h(q)-P(\omega, \vec{x}_{1}(T(q)))\right)} d\omega$$

where $h(q) = -\frac{1}{c_2}(\sqrt{q^2 + x^2} - \sqrt{T(q)^2 + x^2})$. So there are two effects on f in the zero-th order approximation. The first is a phase shift by h(q), under which both the AVFT and the ST are invariant. The second is a small amplitude modulation (for $|\theta_1 - \theta_0| = \pi/6$, this ranges from .94 to 1.077). However, since the error in this approximation is $(\frac{1}{kx})^{1/2} \approx .24$, we should only roughly expect this to hold, since most of the signals have amplitudes on the order of .001. In the full case of Eq. (5), we have a ratio of Bessel functions $A(q,\omega)e^{i\kappa(q,\omega)} := \frac{J_0(k_2\sqrt{q^2+x^2})}{J_0(k_2\sqrt{T(q)^2+x^2})}$ whose argument depends linearly on the frequency ω . If $A(q,\omega) \leq 1$, then this can definitely be written in the form of a non-constant frequency shift $\omega(t)$ as in Theorem 1.

4.3 Effects of translation

Increasing the distance x to the object from x_1 to x_2 is a similar transformation to rotation, since every point on the new rail corresponds to a point on the old rail, but with increased radius. If r is the location on the new rail then the point with the same angle is just $T(r) = \frac{x_1}{x_2}r$; since $x_2 > x_1$, this is smaller than y. Then we have the same sort of derivation as above, with a strictly linear function instead of Eq. (4).

For translation along a given rail, the dependence on $\Theta(\theta)$ is unavoidable. For a given point $x_{\theta}(r)$, the radius is simply $\sqrt{r^2 + x^2}$, while the angle is $\varphi(r) = \theta + \arctan(r/x)$. We leave further discussion of the properties of Θ for our future publication.

5. RESULTS

As noted before, for the linear classifier applied after the non-linear transform, we used sparse logistic regression as implemented in **GLMNET**.¹⁹ For the scattering transform, we used the **SCATNET** package as implemented by Mallat and his group.⁷

For the synthetic dataset, we compared three transforms. The first was the absolute value of the Fourier transform (AVFT). The second, hereafter referred to as the coarser ST, was a three layer scattering transform using Morlet Wavelets, with different quality factors rates in each layer: $Q_1 = 8, Q_2 = 8, Q_3 = 1$. Increasing quality factor corresponds to decreasing the rate of scaling of the mother wavelet, and thus gives more coefficients for the same frequency regime. The third, hereafter referred to as the finer ST, was another three layer scattering transform with slightly lower quality in the second layer, and increased quality in the third layer: $Q_1 = 8, Q_2 = 4, Q_3 = 4$. For each of these, we used 10-fold cross validation to check the generalization of our results.

As in section 4, we investigated two problems: shape and material discrimination. For the shape discrimination, we compared the triangle and the shark-fin, with the material speed of sound fixed at $c_1 = 2000$ m/s, while for the material discrimination, we fix Ω to be a triangle, and compared $c_1 = 2000$ m/s with $c_1 = 2500$ m/s. To do this comparison, we use the received operator characteristic (ROC) curve, which compares the trade off between false positives and true positives as we change the classification threshold; since it strictly concerns one class, it is insensitive to skewed class sizes.²⁴ A way of summarizing the ROC is the area under the curve (AUC), which simply integrates the total area underneath the curve; we use the trapezoidal approximation. This varies from .5 for random guessing to 1 for the ideal classifier, which doesn't misclassify.

The results for material discrimination are in Fig. 8; the corresponding AUCs are .992837 for the AVFT, .97778 for the coarser ST, and .99994 for the finer ST. Fitting with the basic derivation in the geometry section, even



Figure 8: The ROC curve for detecting the material difference in a triangle, for speeds of sound $c_1 = 2000$ m/s and $c_1 = 2500$ m/s. Note that the finer ST curve is an ideal classifier, completely in the upper left. The diagonal line is equivalent to random guessing.

Table 1: The objects in each class for the real dataset

· · · · · · · · · · · · · · · · · · ·	
UXO-like	Other Objects
155mm Howitzer with collar	55-gallon drum, filled with water
$152 \mathrm{mm} \mathrm{TP}-\mathrm{T}$	rock
155mm Howitzer w/o collar	2ft aluminum pipe
aluminum UXO replica	Scuba tank, water filled
steel UXO replica	
small Bullet	
DEU trainer (mine-like object)	

the AVFT is capable of discriminating material effectively. Somewhat surprisingly, the coarser ST performed worse than the AVFT. This is likely because of insufficient frequency resolution, which the finer ST is able to achieve.

The results in shape discrimination are more definitive in demonstrating the effectiveness of the ST, as seen in Fig. 9. The coarser scattering transform, with an AUC of .886, outperforms the AVFT with an AUC of .775. But the finer ST clearly outperforms both of these, with an AUC of .998, on par with the classification rates for the speed of sound problem.

For the real dataset, we compared two transforms, the AVFT and a two layer scattering transform with $Q_1 = Q_2 = 8$. We have split the data into a set of objects that are either UXOs or replicas, and a set of the other objects in the dataset, as listed in Table 1. In both classes, there are a variety of materials and shapes. Between classes, there are no similar shapes (as the shape is what determines if it is a replica rather than a UXO), but there are two with the same material (aluminum UXO replica vs aluminum pipe). The ST has an AUC of .9487, while the AVFT has an AUC of .8186. The AVFT actually did better on this problem than it had on the shape detection problem, suggesting that it is primarily the material properties of the UXOs that distinguish them.



Figure 9: The ROC curve for discriminating a shark-fin from a triangle where both have a speed of sound fixed at 2000m/s.



Figure 10: The ROC curve for detecting UXOs.

6. SUMMARY

In this paper, we have given initial results on understanding the problem of classifying sonar signals using the scattering transform. Using geometric arguments and the synthetic dataset, we have demonstrated that material detection is a considerably simpler problem than shape detection, and that the scattering transform is capable of solving both problems. We have given arguments for why the scattering transform may be suited to this problem by using previous theory, though it is difficult to accurately cast the transform in Eq. (5) in the form of a non-linear frequency modulation. Further, we have given numerical evidence that the scattering transform works

well on this problem. This suggests that there may be additional operators for which the scattering transform has the Lipschitz stability of Eq. (2), which better characterize the object domain.

ACKNOWLEDGMENTS

This research was supported in part by the grants from ONR N00014-12-1-0177, N00014-16-1-2255, and from NSF DMS-1418779. We would like to thank Frank Crosby and Julia Gazagnaire of Naval Surface Warfare Center, Panama City, FL, for providing us with the real BAYEX14 dataset. The work on the synthetic dataset is extended from the work of our former intern Vincent Bodin (now at Sinequa) and our former postdoctoral researcher Ian Sammis (now at Google) who first generated the dataset. James Bremer (UC Davis) and Ian Sammis developed the fast Helmholtz equation solver that we used. Finally, we would like to thank Stephane Mallat and his group (ENS, France) for their ST codes, Simon Kornblith (MIT) for the Julia wrapper of the GLMNET Fortran code, which was in turn developed by Jerome Friedman, Trevor Hastie, Rob Tibshirani and Noah Simon (Stanford).

REFERENCES

- [1] LeCun, Y., Bengio, Y., and Hinton, G., "Deep learning," *Nature* **521**(7553), 436–444 (2015).
- [2] Krizhevsky, A., Sutskever, I., and Hinton, G. E., "ImageNet classification with deep convolutional neural networks," NIPS 25, 1097–1105 (2012).
- [3] Sun, Y., Wang, X., and Tang, X., "Deep convolutional network cascade for facial point detection," Proc. IEEE Trans. Pattern Anal. Mach. Intell., 3476–3483 (2013).
- [4] Dahl, G. E., Sainath, T. N., and Hinton, G. E., "Improving deep neural networks for LVCSR using rectified linear units and dropout," *IEEE Int. Conf. on Acoustics, Speech and Signal Process.*, 8609–8613 (2013).
- [5] Mallat, S., "Understanding deep convolutional networks," *Philos trans. A, Math. Phys eng sci* 374 (2016).
- [6] Mallat, S., "Group invariant scattering," Comm. Pure Appl. Math. 65(10), 1331–1398 (2012).
- [7] Bruna, J. and Mallat, S., "Classification with scattering operators," IEEE Conf. Comp. Vision and Pattern Recog., 1561–1566 (2011).
- [8] Wiatowski, T. and Bölcskei, H., "Deep convolutional neural networks based on semi-discrete frames," *IEEE Int. Symp. on Info. Theory*, 1212–1216 (2015).
- [9] Wiatowski, T., Tschannen, M., Stanic, A., Grohs, P., and Bölcskei, H., "Discrete deep feature extraction: A theory and new architectures," *Proc. Int. Conf. on Machine Learning* (2016).
- [10] Chen, X., Cheng, X., and Mallat, S., "Unsupervised deep Haar scattering on graphs," Adv. Neural Info. Process Syst 27, 1709–1717 (2014).
- [11] Mallat, S. and Waldspurger, I., "Phase retrieval for the cauchy wavelet transform," J. Fourier Anal. Appl. 21(6), 1251–1309 (2015).
- [12] Chudáček, V., Andén, J., Mallat, S., Abry, P., and Doret, M., "Scattering transform for intrapartum fetal heart rate variability fractal analysis: a case-control study," *IEEE Trans. on Biomed. Eng.* 61(4), 1100–1108 (2014).
- [13] Chang, K.-Y. and Chen, C.-S., "A learning framework for age rank estimation based on face images with scattering transform," *IEEE Trans. on Image Process.* 24(3), 785–798 (2015).
- [14] Dov, D. and Cohen, I., "Voice activity detection in presence of transients using the scattering transform," *IEEE Conv. of Electrical Electronics Eng. in Israel*, 1–5 (2014).
- [15] Kargl, S., "Acoustic response of underwater munitions near a sediment interface: Measurement model comparisons and classification schemes," Tech. Rep. MR-2231, SERDP and University of Washington Seattle Applied Physics Lab (2015).
- [16] Marchand, B., Saito, N., and Xiao, H., "Classification of objects in synthetic aperture sonar images," *IEEE/SP Workshop on Stat. Signal Process.* 14, 433–437 (2007).
- [17] Bremer, J., "A fast direct solver for the integral equations of scattering theory on planar curves with corners," J. Comput. Phys. 231(4), 1879–1899 (2012).
- [18] Bremer, J. and Gimbutas, Z., "On the numerical evaluation of the singular integrals of scattering theory," J. Comput. Phys. 251, 327–343 (2013).

- [19] Hastie, Trevor; Tibshirani, R. and Wainwright, M., [Statistical Learning with Sparsity The Lasso and Generalizations], Chapman & Hall/CRC Press (2015).
- [20] Friedman, J., Hastie, T., and Tibshirani, R., "Regularization paths for generalized linear models via coordinate descent," J Stat. Softw. 33(1), 1–22 (2010).
- [21] Otsu, O., "An invariant theory of linear functionals as linear feature extractors," Bulletin of the Electrotechnical Laboratory **37**(10), 893–913 (1973).
- [22] Amari, S., "Invariant structures of signal and feature space in pattern recognition problems," RAAG Memoirs 4(1-2), 553–566 (1968).
- [23] Etter, P. C., [Underwater Acoustic Modeling and Simulation], CRC Press, Boca Raton, FL (2013).
- [24] Fawcett, T., "An introduction to ROC analysis," Pattern Recogn. Lett. 27(8), 861–874 (2006).