

Automated Classification of Food Products Using 2D Low-Field NMR

Mason Greer^a, Cheng Chen^a, Soumyajit Mandal^{a,*}

^a*Case Western Reserve University, 10900 Euclid Ave, Cleveland, OH 44106, USA*

Abstract

In this work, low-field proton (^1H) and sodium (^{23}Na) relaxation and diffusion measurements are used to detect and classify different types of food products. A compact and low-cost system based on a small 0.5 T permanent magnet has been developed to autonomously authenticate such products. The system uses a simple but efficient double-tuned matching network suitable for $^1\text{H}/^{23}\text{Na}$ NMR. Various machine learning algorithms are used to classify food samples based on T_1 - T_2 and D - T_2 data generated by the system, and the accuracy and prediction speed of these algorithms are studied in detail. The influence of temperature drift upon prediction accuracy is also studied. Experimental results demonstrate reliable classification of cooking oils, milk, and soy sauces.

Keywords:

$^1\text{H}/^{23}\text{Na}$ NMR, Double-tuned network, Machine learning, Food products authentication

1. Introduction

Quality control of consumables (food items, dietary supplements, and essential medicines) is of increasing importance in today's society [1, 2, 3, 4]. In particular, finding a relatively inexpensive and practical solution for accurate quality control can help to detect false labeling, adulteration, as well as spoiling of food products. This will allow consumers to know that they are getting the product they think they are buying and give them confidence that it is safe to consume. One of the first steps towards accurate quality control is the accurate and automated classification of food products into categories and sub-categories.

One category of food products of interest is cooking oils because higher cost oils are often cut with cheaper ones to increase profits [5]. Many experimental methods have been used to classify and detect adulterated cooking oil products including near infrared (NIR), mid-infrared (MIR), Raman, and NMR spectroscopy [6]. High-resolution NMR spectra have been previously used to successfully evaluate adulterated cooking oils [7]. However, such NMR techniques require large, complicated, and costly equipment (e.g., superconducting magnets) as well as highly trained personnel to analyze the results. As an alternative, it has been shown that inexpensive and compact low-field NMR systems can accurately classify cooking oils using T_2 relaxation times [8]. However, the classification process was not automated.

Soy sauces are another food product prone to adulteration and mislabeling in order to increase profits for manufacturers. Soy sauce can be produced in two separate methods. The first method involves fermentation by microorganisms and is known as fermented soy sauce (FSS), while the second uses acids to quickly digest the proteins and carbohydrates and is known as blended soy sauce (BSS) [9, 10]. In some cases, manufacturers have intentionally mislabeled BSS packages as FSS packages in order to improve profits [10]. In other cases, the soy sauce itself is fake and is produced by mixing brine, monosodium glutamate, and dark soy sauce [11]. Current methods for evaluating the authenticity of soy sauces include NIR spectroscopy, capillary electrophoresis, and liquid chromatography [11, 10, 12]. ^1H and ^{13}C NMR have been used to study the compositional differences between different types of soy sauces [13, 14]. However, we are not aware of any study using low-field ^1H or ^{23}Na to study and classify different types of soy sauces. Because of this, soy sauce authentication faces the same constraints that cooking oils face.

Adulteration and spoiling of milk is another area of large concern. Currently, retailers use printed expiration dates on milk products to detect spoiling. However, this method is largely inaccurate, resulting in unspoiled milk being discarded or spoiled milk being sold [15]. The former results in a loss of money for retailers and a loss of perfectly good food product. The latter can potentially lead to food poisoning to the consumer. Current methods of detection for the spoilage of milk are pH testing, electrical tests using an amperometric sensor, magnetoelastic sensing, gas-sensor arrays, IR spectroscopy, and protein/fat count detection

*Corresponding author

Email addresses: mkg53@case.edu (Mason Greer), cxc717@case.edu (Cheng Chen), sxm833@case.edu (Soumyajit Mandal)

[16, 17, 18, 19, 20, 21]. The use of low-field NMR to detect spoiled milk is a relatively unexplored area.

In this paper, we address the size, cost, and operator constraints of NMR for classification of food items by introducing a dual-tuned, low-field NMR system that has the potential to autonomously classify different food products and/or detect spoilage. We test the feasibility of the device through collecting ^1H T_1 and T_2 relaxation data on various cooking oils. We also study the T_1 , T_2 , and diffusion of soy sauces with low-field ^{23}Na NMR. Finally, we use T_1 - T_2 maps to characterize fresh versus spoiled milk. The data collected is preprocessed and used to train a variety of machine learning algorithms for automatic classification of each sample. Finally, we compare the accuracy of each of the models to each other. Our work extends earlier results on NMR-based analysis of food products [22] by focusing on i) 2D low-field ^1H and ^{23}Na NMR, and ii) modern machine learning and classification algorithms.

2. System Overview and Theory

A block diagram describing the proposed automated system is shown in Figure 1. The system operates as follows. The user first selects the food product they will be scanning. Based on the product, the system selects appropriate pulse sequences, machine learning algorithms, and trained models from a database. The number of scans is set to obtain sufficient signal-to-noise ratio (SNR) as set by the user. The acquired data is collected from the scanner and processed accordingly. To extract the 2D T_1 - T_2 and the D - T_2 correlation maps, the inverse 2D Laplace transform is used following the procedure detailed in [23]. After the data processing, key features are extracted in order to classify the sample. In this paper, we first set a threshold in order to help remove possible fitting errors. Gaussian functions are fitted to each

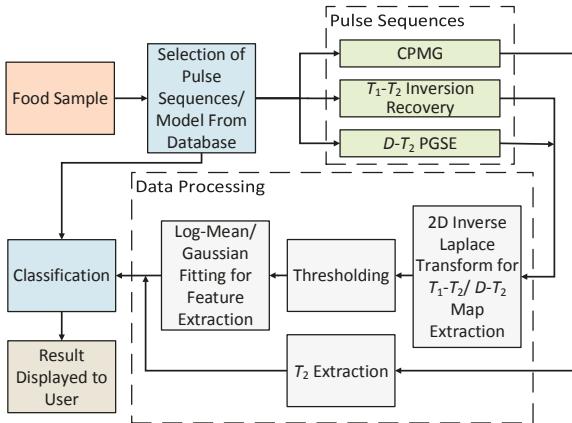


Figure 1: A flowchart of the proposed NMR-based method for automated classification of food products.

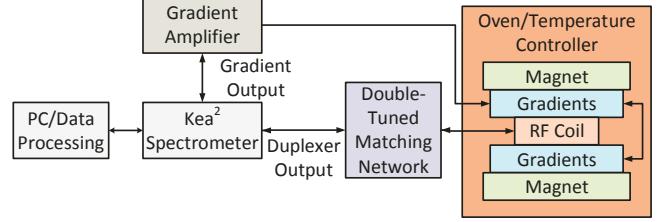
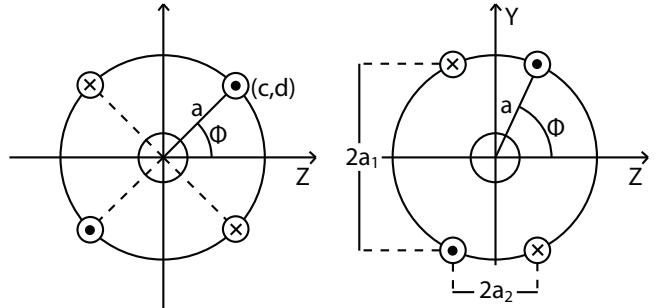


Figure 2: Block diagram of the proposed low-field NMR system for automated classification of food products.



(a) The ideal z-gradient geometry where $\phi = 45^\circ$. (b) The distorted z-gradient geometry where $\phi \neq 45^\circ$.

Figure 3: Geometry of the gradient coils used in the proposed system.

axis of the map to get information about the T_1 , T_2 , and diffusion coefficient D of the sample. Alternatively, log-mean values can be used to extract similar information in cases of low SNR. The covariance matrix between the two axes of the maps can also be extracted to retrieve information about the shape of the map.

The NMR system consists of a 0.5 T desktop permanent magnet (SpinCore, Gainesville, FL) for generating the B_0 field. We use one of two custom homebuilt probes with y and z gradients that are capable of holding either 5 mm or 10 mm NMR tubes. The gradients can be used for both 2D imaging and diffusion measurements but are used for diffusion measurements in this paper. Because the system needs to be able to collect both ^1H and ^{23}Na NMR data, a double-tuned matching network is needed. A desktop spectrometer (Kea2, Magritek) is used for RF and gradient pulse control and data collection but a custom spectrometer board designed in our lab [24] can be used instead to reduce the size and cost of the system. Commercial gradient amplifiers (AE Techron 7224) are used to drive the gradient coils, but custom miniaturized gradient driver boards have also been developed to further reduce system size. A block diagram of the entire system is shown in Figure 2.

2.1. Gradient Coil Design

Each gradient coil is constructed using four parallel wires. The ideal cross-section of the gradient geometry is shown in Figure 3a, in which each wire is located 90° apart from

each other along a circle of radius a . In practice, this ideal geometry may not be realizable due to probe design constraints, particularly if the top and bottom halves of the coil are fabricated on two printed circuit boards (PCBs) for convenience. Thus, the geometry of the gradients may need to become slightly distorted in order to satisfy these constraints, as shown in Figure 3b. Here the wires are still positioned symmetrically about the y and z axes, but now have arbitrary vertical and horizontal separations of $2a_1$ and $2a_2$, respectively.

The ideal geometry can be solved analytically [25]. In Appendix A, we show this derivation and then use it to derive the field and gradient for the distorted case. Summarizing the results found in Appendix A, we get gradients of

$$\begin{aligned} g_y &\equiv \frac{\delta B_z}{\delta y} = \left(\frac{2\mu_0 I}{\pi a^2} \right) \cos(2\phi) \\ g_z &\equiv \frac{\delta B_z}{\delta z} = \left(\frac{2\mu_0 I}{\pi a^2} \right) \sin(2\phi) \end{aligned} \quad (1)$$

for the ideal geometry, where I is the current in each wire. Eqn. (1) shows that g_z will reach its maximum value when $\phi = \pi/4$. This will also make the y -gradient go to 0, leaving only a z -gradient. The opposite can be said if all four wires are rotated by 45° , i.e. $\phi = 0$. Thus, the y -gradient can be generated by another set of four wires with $\phi = 0$.

As stated above, the ideal gradient geometry is not always feasible in practice. The distorted geometry shown in Figure 3b can be used instead, at the expense of some gradient strength and linearity across the sample. The gradient strength of the distorted geometry is shown below:

$$\begin{aligned} g_y &\equiv \frac{\delta B_z}{\delta y} \approx \left(\frac{2\mu_0 I}{\pi a^2} \right) \left[6 \sin(4\phi) \frac{yz}{a^2} + \dots \right], \\ g_z &\equiv \frac{\delta B_z}{\delta z} \approx \left(\frac{2\mu_0 I}{\pi a^2} \right) \left[\sin(2\phi) + 3 \sin(4\phi) \frac{(y^2 - z^2)}{a^2} + \dots \right]. \end{aligned} \quad (2)$$

Analyzing the gradients shown in eqn. (2) near the origin, we see that the z gradient is maximum at an angle of $\phi = \pi/4$, which agrees with the original analysis of the optimal geometry. However, the distorted geometry has $\phi = \tan^{-1}(a_2/a_1) \neq 45^\circ$ resulting in lower z -gradient strength by a factor of $\sin(2\phi)$ at the same current. The gradient will also be less uniform over the sample due to the introduction of third-order terms into the magnetic field, which is a consequence of the lower symmetry of the geometry. Moreover, the gradients are coupled, i.e., the g_z cannot be maximized while keeping $g_y = 0$, and vice-versa. Therefore, if a y -gradient is to be added, it should be kept as close to an ideal geometry as possible. This can be accomplished by adding i) a third wire to both existing PCBs located at $(\pm a_1, 0)$, i.e., on the y -axis; and ii) two wires at $(0, \pm a_1)$, i.e., on the z -axis.

2.2. Double-Tuned Matching Network Design

We would like to design a double-tuned impedance-matching network to perform $^1\text{H}/^{23}\text{Na}$ double-resonance experiments with a single RF coil. This would allow us to interleave ^1H and ^{23}Na pulse sequences to reduce overall measurement time. However, the power efficiency of such networks is of concern, since it is generally lower than for single-tuned networks. The efficiency η of any impedance-matching network is defined as the ratio of power delivered to the coil P_{coil} to the total input power P_{in} . Both are functions of frequency, so in general we write the efficiency as $\eta(f) = P_{\text{coil}}(f)/P_{\text{in}}(f)$. An ideal matching network is lossless, resulting in $\eta = 1$. Most of the power lost within practical matching networks is dissipated within inductors, since they usually have much lower quality factors than capacitors. We would like to design a network that provides high efficiency at both the proton and sodium Larmor frequencies. A suitable network for matching to two such widely-separated frequencies is shown in Figure 4 [26]. Unlike most doubly-tuned networks, it has a single input/output port for both frequencies. The circuit is similar to the common two-capacitor matching network, but contains three extra components: two inductors (L_2 and L_3) and one capacitor (C_2). The $L_2 - C_2$ branch is referred to as a “trap” circuit. At the lower matched frequency it is mostly capacitive, resulting in an effective tuning capacitance of $C_1 + C_2$. At the higher frequency it is almost an open circuit, resulting in a smaller tuning capacitance (approximately C_1). The $L_3 - C_3$ branch then matches the impedance of the tuned coil to the desired real value.

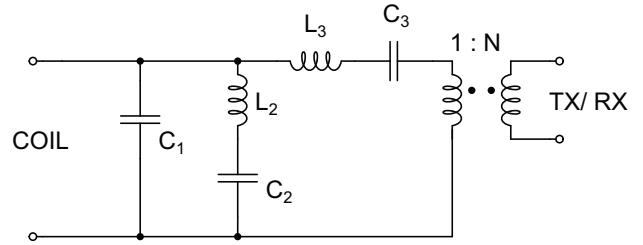


Figure 4: Schematic of the double-tuned tuning and impedance-matching network used in the experiments.

We chose this probe circuit because it uses very few additional passive components while providing fairly high efficiency (low internal power losses) at both matching frequencies. Moreover, the single input/output port allows it to be used with low-cost single-channel spectrometer consoles. The $1:N$ transformer at the output changes the matched input impedance of the rest of the network to $Z_{\text{in}} = Z_0/N^2$, where $Z_0 = 50 \Omega$ is the standard impedance of the cables, connectors, and spectrometer. In our case we used $N = 2$, resulting in $Z_0/N^2 = 12.5 \Omega$. Lowering the impedance reduces the required values of inductors L_2 and L_3 , allowing them to be realized with high- Q air-core inductors of rea-

sonable size. Maximizing Q in turn minimizes power losses within the matching network and increases its efficiency.

The circuit shown in Figure 4 is not unique. Several complementary versions can be created by converting series branches within the circuit into parallel branches, and vice versa [26]. It can also be extended to an arbitrary number of matching frequencies by adding more LC traps. For N frequencies we would need $(N - 1)$ traps.

Another useful property of the matching network shown in Figure 4 is that parasitic reactances can be absorbed into it, resulting in a more robust design. For example, the leakage inductance of the transformer acts as an inductor of value $L_l = L_s(1 - k^2)$ in series with the primary winding, where L_s is the inductance of the primary winding and k is the coupling coefficient between the windings. However, its effects are easily removed by changing the value of L_3 to $L_3 - L_l$. In addition, C_1 can be reduced to account for the inter-turn capacitance of the NMR coil.

Loss within each inductor in the matching network was modeled with a frequency-dependent series resistor of value $R_i = \omega L_i/Q$, where we assumed $Q = 70$ and 120 at the sodium and proton Larmor frequencies of 5.8 MHz and 22 MHz, respectively (these values are typical for the Coilcraft 2222SQ and 2929SQ series of surface-mount air-core inductors used in our experiments). Loss within the NMR coil L_1 was modeled similarly, but with the series resistance derived from a polynomial fit to the measured frequency-dependent series resistance. The efficiency η of the network can then be analytically calculated for any choice of component values. We begin by defining η as

$$\eta = \frac{\text{Power loss in NMR coil}}{\text{Total power loss}} = \frac{P_1}{\sum_{i=1}^3 P_i}, \quad (3)$$

where $P_i = |I_i|^2 R_i/2$ is the power dissipated in each coil. We further define $TF_i \equiv I_i/I_{in}$ as the transfer function between current in the input terminal and the i -th coil. These are given by

$$TF_1 = \frac{(Z_{in} - Z_3)}{sL_1 + R_1}, \quad TF_2 = \frac{(Z_{in} - Z_3)}{sL_2 + R_2 + 1/(sC_2)},$$

and $TF_3 = 1$, (4)

where $s \equiv i\omega$ is the complex frequency variable. Finally, $\eta(\omega)$ is given by

$$\eta(\omega) = \frac{|I_1|^2 R_1}{\sum_{i=1}^3 |I_i|^2 R_i} = \frac{|TF_1|^2 R_1}{\sum_{i=1}^3 |TF_i|^2 R_i}. \quad (5)$$

Losses within the transformer have been neglected in this analysis, but can be easily included if necessary. However, it is difficult to analytically calculate the correct component

values for this network, because of its complexity. In addition, network synthesis is usually a many-to-one problem, meaning that many different combinations of component values will provide similar or identical performance characteristics. Hence we used numerical optimization to synthesize the network. The cost function that was minimized during the optimization was a weighted sum of the reflection coefficient $\Gamma \equiv |Z_{in} - Z_0| / |Z_{in} + Z_0|$ and power loss within the network at both design frequencies:

$$C = \sum_{i=1}^2 |\Gamma(f_i)| + \alpha \sum_{i=1}^2 [1 - \eta(f_i)], \quad (6)$$

where the parameter α can be adjusted to trade-off between impedance matching and efficiency. Such flexibility resulted in better solutions than directly maximizing P_1 (power delivered to the RF coil), which is given by $P_{in} (1 - |\Gamma(f_i)|^2) \eta(f_i)$ at each matching frequency f_i .

Nevertheless, we found the optimization space to be highly non-convex. As a result, gradient-based optimization algorithms quickly got stuck in local minima, and repeated runs with random initial conditions still produced poor results. We therefore used a stochastic global optimization algorithm for this problem, namely the genetic algorithm (GA) available in the MATLAB Global Optimization Toolbox [27]. Many papers have reported on the application of GAs in optimizing power amplifiers and other active circuits, but their use in designing this kind of matching network appears to be novel.

2.3. Pulse Sequences and 2D Maps

We utilize the following pulse sequences to classify food products: Carr-Purcell-Meiboom-Gill (CPMG) for T_2 measurements, inversion Recovery (IR) with CPMG pulse train for T_1 - T_2 measurements, and pulsed gradient spin echo (PGSE) with CPMG pulse train for D - T_2 measurements. All sequences utilize rectangular RF pulses. The IR and PGSE sequences, which are shown in Figure 5, use CPMG pulse trains to collect T_2 information along with the T_1 or D information provided by the IR or PGSE encoding periods. These correlation experiments thus allow us to obtain 2D probability distributions on the T_1 - T_2 and D - T_2 planes. We can then fit curves to these distributions to extract features for training classification models.

In order to calculate the correlation plots described above, we use the 2D inverse laplace transform (ILT). The algorithm used to efficiently calculate the ILT utilizes the singular-value decomposition (SVD) for compression, and Tikhonov regularization to improve numerical stability for noisy data [23]. Specifically, we minimize the cost function

$$C_{ILT} = \|M - K_1 F K_2'\|^2 + \alpha \|F\|^2, \quad (7)$$

where the first term represents the error between the data matrix M and the fit $K_1 F K_2'$, and the second term helps

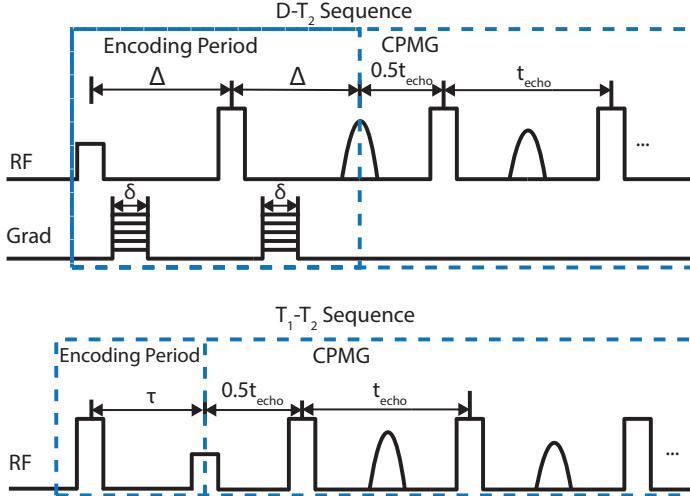


Figure 5: (Top) The $D\text{-}T_2$ correlation experiment. The gradient is stepped over multiple experiments to extract the diffusion data and the decay due to diffusion will be a function of the gradient strength and the delays δ and Δ (Bottom) The $T_1\text{-}T_2$ correlation experiment. The delay τ is stepped over multiple experiments to extract the T_1 data.

to smooth the results (Tikhonov regularization). Also K_1 and K_2 are the measurement kernels, and F is the desired 2-D probability distribution. The value of α determines the amount of regularization, i.e., smoothing of F . However, large amounts of regularization bias the fit results. The optimal value of α can be chosen by plotting the fitting error versus α . As α is increased, the overall fitting error will remain nearly constant before starting to sharply increase. We chose the largest α for which the error remains close to its minimum value. This produces the smoothest possible probability distribution F that does not significantly worsen the bias. More formal methods for optimizing α given the data and the noise variances, such as the Butler-Reeds-Dawson method [28], can also be used.

2.4. Machine Learning and Classification

Recent advancements in machine learning suggest its importance in interpreting and classifying complex data-sets for decision-making. In particular, recent work on discrimination and authentication of food items has relied on state-of-art machine learning algorithms [29, 30, 31]. In this paper, we use the inverted 2D data sets to extract parameters for classification. It should be noted that classification can also be performed on the raw data. However, since the latter are very large (typically $\approx 5 \times 10^6$ points each), classifying them would be slow and computationally intensive. We therefore greatly reduce the sizes of our data sets (typically, to 15×30 matrices) by using SVD. The compressed data matrices are suitable for fast classification, but provide little insight into sample properties. We thus prefer to invert the compressed matrices and then perform classification on

the resulting probability distributions, which provides more physical insight.

Many analytical chemistry classification studies based on NMR or other spectroscopic techniques implement classical chemometric classification methods such as principal component analysis (PCA) or partial least squares discriminant analysis (PLS-DA) [14]. The aim of PCA is to reduce the dimensionality of a data set while capturing as much of the original variance as possible. This works well for expressing the data set in terms of a smaller number of linearly uncorrelated variables (the principal components), but does not by itself enable automated classification of an unknown sample. Thus, PCA is mainly useful as a tool for feature selection, i.e., improving the prediction accuracy of classification algorithms by reducing the dimensionality of the problem. However, using PCA as a feature selection tool *after* inversion does not work well because the variance of the inverted data is sensitive to noise and the value of the regularization parameter α . Thus, we do not include PCA-based feature selection models in our results.

PLS-DA works well for high-dimensional data sets such as a full NMR spectrum but does not work as well in a low-dimensional space [32]. In our case, we reduce the dimensionality of our data by extracting specific T_1 , T_2 , or diffusion coefficient values from their distributions, making PLS-DA a poor choice. Because of this, we choose to avoid some of the more classical chemometric methods and move towards other methods used commonly for classification. In particular, we utilize the MATLAB Classification Learner application to test a variety of modern machine learning algorithms on the acquired data. We then use cross-validation (as opposed to other methods such as bootstrapping) to test the accuracy of each algorithm. This is because studies have shown that cross validation generally outperforms bootstrapping for the same amount of computation [33].

In particular, we focus on four different categories of machine learning algorithms readily available in the Classification Learner application. These categories are decision trees, discriminants, support vector machines (SVM), and k -nearest neighbors (kNN). A brief description of these algorithms is given below; a more detailed description can be found in [34].

Decision trees classify data based on specific feature values. They work by building a large tree of nodes that separate the data by comparing the specific features to a previously-built model. As a data set moves through more nodes, it will eventually become classifiable based on the path the data point follows through the decision tree. Decision trees are often used because they are easily comprehensible.

Discriminant analysis methods work in a similar fashion to PCA, since both attempt to reduce data dimensionality. However, while PCA reduces dimensionality by selecting components that best explain the observed variance in

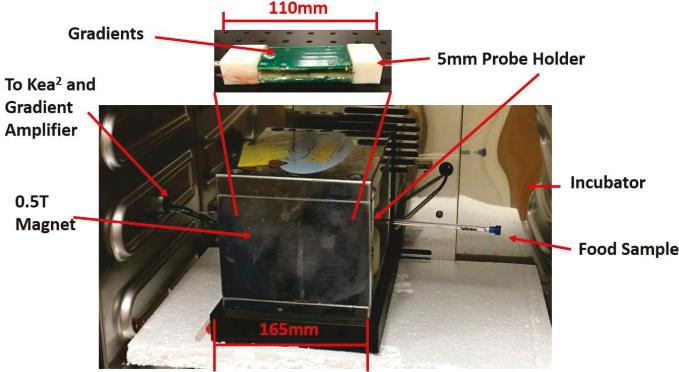


Figure 6: The measurement setup. A home-made 3D-printed probe holder is inserted in the 0.5 T permanent magnet. Gradient circuits are mounted on each side of the holder. The figure shows an example of the probe holder for the 5 mm tube. The setup is placed in an incubator to reduce the influence of temperature drifts.

the data, discriminant methods reduce dimensionality by selecting data labels that maximize the separation between known classes in the training set.

SVMs try to separate the data by drawing a hyperplane between the labeled training set categories. When drawing the line, the algorithm tries to maximize the distance or margin between the different categories. When testing a new data point, the algorithm checks to see what side of the hyperplane the point is on. Based on this information, it can classify a data set. For linearly separable data, the hyperplane can be found relatively easily and data can be classified. In other instances, the data may not be linearly separable, but separable in a different higher dimensional space. In this instance, a kernel function can be used to map the data into a higher dimensional space, which can then be used for classification. In general, SVM is a binary classifier, which means that it can only classify data into two classes. However, multiple SVM models can be combined together to handle any number of classes.

kNNs work under the assumption that data classified under the same label will have features that lie near each other when plotted. Thus, the algorithm compares a labeled training data set to a test data set that is to be classified by looking for the k nearest neighbors to the test data set, where k is an integer. It then classifies the data set by finding the most frequently used label among the nearest neighbors. *kNN* has been shown to work well in many instances, but a major drawback is its relatively high computational cost.

3. Results

The overall experimental setup is shown in Fig. 6. The following sections describe the hardware and experimental results in more detail.

3.1. Gradient Coils

The gradient coils were implemented on four separate two-layer PCBs with thick (2 ounce) copper traces to reduce resistance and a thickness of 0.8 mm. Parallel wires were placed on both sides of the PCB to increase the gradient strength. The same PCBs were used with two probes: one for 10 mm NMR tubes and another for 5 mm tubes. The boards were mounted on custom-designed 3D-printed probe holders made of ABS (Acrylonitrile Butadiene Styrene) plastic, which is strong enough to hold the gradients, RF coil, and sample. Figure 6 shows one of the assembled probe holders. The 10 mm gradient geometry is more suitable for imaging as opposed to diffusion measurements due to its smaller maximum gradient strength, but is acceptable for samples with relatively long relaxation times.

In the 10 mm system, the z -gradient PCBs are separated by 15 mm in the z -direction, while traces on the PCBs are separated by 12 mm along the y -direction. Because the PCBs have finite thickness, the wires along the z -direction have different a_1 values depending on whether they are on the top or bottom of the board. Thus, we calculate the fields separately from each layer and then use superposition. The resulting values are listed in Table 1. The y -gradients follow the ideal geometry in this configuration allowing us to use eqn. (1). The resulting values are listed in Table 2.

We choose to use the y -gradient and the 5 mm probe for the experiments as they provide the strongest and most uniform gradient. The gradient strength was calibrated using PGSE sequences on a known sample (water). The result was 1.84 T/m at 30 A, which is 10% larger than the expected value. This discrepancy is probably due to assembly tolerances, i.e., the PCBs being slightly closer than expected.

We compare the linearity of our distorted geometry for the 5 mm probe with an ideal geometry to make sure that the probe is suitable for both diffusion measurements and imaging. Figure 7 shows the computed fields and gradients for both geometries versus distance. The ideal geometry has a somewhat larger and more uniform gradient, as expected. However, both have acceptable performance across the sample, i.e., for $|z| < 2.5$ mm.

3.2. Double-Tuned Matching Network

Repeated GA runs turned up a number of very good solutions, which are shown in Figure 8. The plot on the top shows that the power efficiencies at the two frequencies are inversely correlated, such that their average η_{av} remains approximately constant. This behavior is an intrinsic property of this matching network topology [26]. The optimization procedure outlined above can be easily generalized to other matching problems. For example, we have successfully used it to generate matching networks for $^1\text{H}/^{13}\text{C}$ operation with the same NMR coil [35].

We chose solution 15 for implementation because it has high power efficiency and also a relatively small value of L_3 ,

Probe size	$a_{1_{in}}$	$a_{2_{in}}$	$a_{1_{out}}$	$a_{2_{out}}$	Radius _{in}	Radius _{out}	ϕ_{in}	ϕ_{out}	Calculated g_z
10 mm	6 mm	7.5 mm	6 mm	8.3 mm	9.6 mm	10.1 mm	38.7°	35.3°	15.7 mT/m/A
5 mm	6 mm	5 mm	6 mm	5.8 mm	7.8 mm	8.3 mm	50.2°	46.0°	24.3 mT/m/A

Table 1: The z -gradient geometry parameters and their corresponding calculated gradient strengths.

Probe size	Radius _{in}	Radius _{out}	ϕ_{in}	ϕ_{out}	Calculated g_y
10 mm	7.5 mm	8.3 mm	0°	0°	26.0 mT/m/A
5 mm	5 mm	5.8 mm	0°	0°	55.7 mT/m/A

Table 2: The y -gradient geometry parameters and their corresponding calculated gradient strengths.

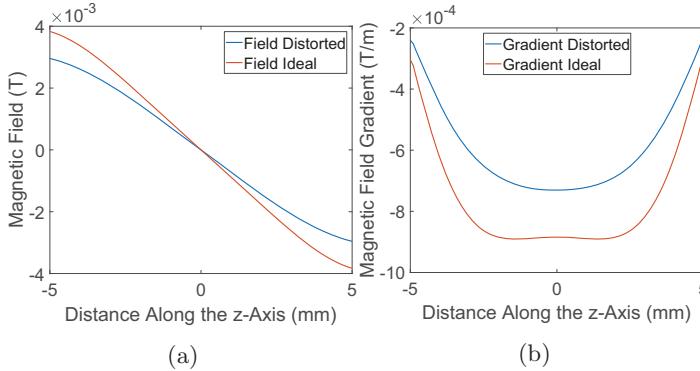


Figure 7: A comparison of the magnetic field and the magnetic field gradient generated by the ideal gradient geometry and the distorted geometry. (a) Magnetic field profiles, and (b) gradient profiles.

allowing it to be implemented with two air-core inductors (Coilcraft 2222SQ or 2929SQ) in series. The matching network was implemented on a two-layer PCB. The bottom layer was generally used as a ground plane. However, the lack of a high- μ core means that the air-core inductors L_2 and L_3 have significant fringe fields. These can cause extra losses by generating eddy currents on nearby metal surfaces. We therefore removed the ground plane near these inductors. A high-performance surface-mount RF transformer (Coilcraft PWB-4-AL, 3 dB bandwidth 0.15 - 500 MHz) was used to match the output to 50Ω . The capacitors were implemented using fixed surface-mount high- Q mica components (Cornell Dubilier) in parallel with mechanically-tunable variable components (Sprague-Goodman, 1 - 20 pF). The power handling capability of the matching network is determined either by the maximum safe operating voltage of the tunable capacitors (100 V), or by the maximum current within the transformer that does not saturate its ferromagnetic core (~ 250 mA); this limit is usually reached first.

Figure 9 shows the measured reflection coefficient $|\Gamma|$ of the assembled coil and double-tuned matching network before final trimming. Good matching is evident at frequencies around 22.03 MHz and 5.83 MHz, which are very close

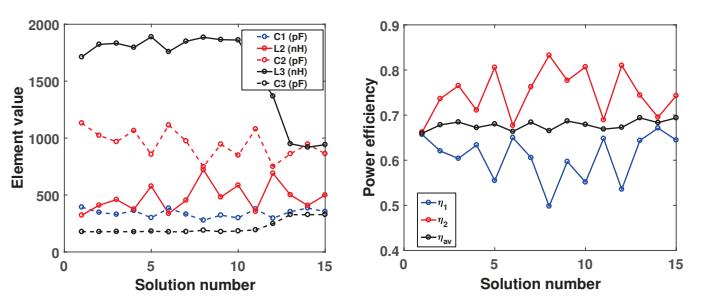


Figure 8: Some of the best solutions found by the GA-based matching network optimization algorithm. Synthesized element values are shown on the left, and power efficiency at both matching frequencies on the right. The average power efficiency over both frequencies is also shown.

to the desired values for this magnet. Since a significant fraction of the input power is dissipated within the inductors L_2 and L_3 , efficiency cannot be directly estimated from Γ . Instead we have to compare the B_1 amplitude produced by the double-tuned network for both ^1H and ^{23}Na with that produced by the single-tuned ^1H and ^{23}Na networks for the same input power levels. The simulation predicts 74% and 64% power efficiency at 22.03 MHz and 5.83 MHz respectively, while the measurement shows 53% and 50% at these two frequencies. The observed reduction in efficiency is probably due to additional loss mechanisms that were ignored in the simulations, such as i) finite capacitor Q , and ii) core and winding losses within the transformer.

3.3. Sample Preparation

All food products tested in the following sections were purchased at a local grocery store. Samples for training and testing were drawn from the product bottles with a syringe and injected into NMR tubes prior to data collection. The products were homogeneous and our tests were intended as a feasibility study, so data was collected and analyzed from a single sample of each product.

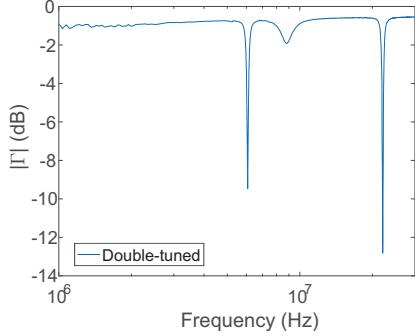


Figure 9: Measured reflection coefficient of the double-tuned probe before final trimming. Minima in the reflection coefficient occur at 22.05 MHz (-9.5 dB) and 5.84 MHz (-12.8 dB). These frequencies are already very close to the desired values of 22.03 MHz for ^1H and 5.83 MHz for ^{23}Na .

3.4. Cooking Oil Experiments

We tested the feasibility of classifying common cooking oils (olive, canola, vegetable, and corn oil) using the proposed setup. For this purpose 100 ^1H T_1 - T_2 correlation maps were collected for each sample. Table 3 shows the experimental parameters used for this experiment. Averaging is not needed as $\text{SNR} \approx 2000/\text{scan}$ for these experiments. Typical maps collected for each sample are shown in Figure 10. The T_1 - T_2 distribution is spread along the diagonal ($T_1 = T_2$ line) in a relatively similar fashion for all the samples. Thus, the qualitative T_1 - T_2 distributions of the samples are similar. However, we show that enough information can be extracted from each distribution for accurate classification.

We fit both the T_1 and T_2 axes to bi-Gaussian functions and extract the two peak locations for each oil as classification features. The covariance matrix between the T_1 and T_2 fits can also be used for classification. We plot four different combinations of the classification features in Figure 11, where a and b reference the first and second components of the Gaussian fit, respectively. All the plots in Figure 11 show clear separation between the samples, suggesting that accurate classification is possible.

The MATLAB Classification Learner app was used to train and test several machine learning algorithms using the features selected above. We used 50-fold cross validation to protect against over-fitting. A summary of the results obtained from all the algorithms is shown in Figure 12 for models trained on three different sets of features. Note that we show the results from all 20 algorithms that were tested to emphasize that a wide range of algorithms work well and provide similar prediction accuracies. Only the cosine kNN is significantly worse than the others: its distance metric (cosine distance) appears to be unsuitable for our data sets. Thus, for later experiments, we only show classification results from a subset of the algorithms for simplicity; a summary of the results from all 20 algorithms can be found in Appendix B.

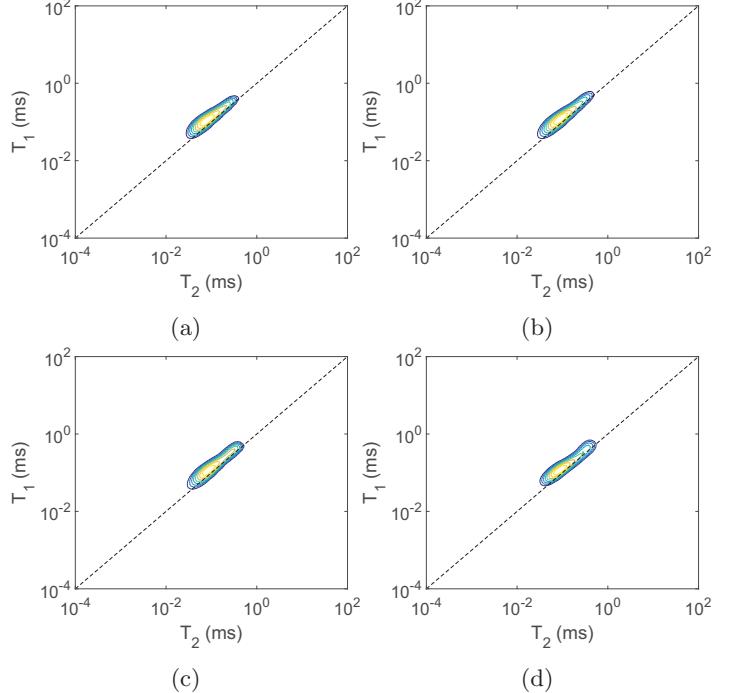


Figure 10: T_1 - T_2 correlation maps collected on (a) olive oil, (b) canola oil, (c) corn oil, and (d) vegetable oil. It is hard to see any qualitative differences between the maps. Each map has $\text{SNR} \approx 2000$ with no averaging applied during the data collection.

All models perform well, with most of the classification accuracies being $> 99\%$. The feature set that includes all the features, including the covariance, performs the worst. This is because covariance values have much more variability between experiments than mean values, which decreases classification accuracy. When the covariance parameters are removed, the performance of most algorithms improves, with some achieving 100% accuracy. The fact that the best results are obtained with only T_{1a} and T_{2b} as features suggests that in general the minimal number of extracted parameters should be retained for classification.

Other important performance metrics of the classification algorithms include training time and prediction speed, which should both be as low as possible. Figures 13 and 14 show these parameters for models trained on the T_{1a} - T_{2b} feature set. The medium tree algorithm is the most efficient algorithm for classifying the cooking oil samples.

3.5. Milk Experiments

We collected T_1 - T_2 correlation maps for a milk sample that was not refrigerated over two days to detect spoiling. The pulse sequence parameters for the milk experiments can be seen in Table 4. Each data set has an SNR of approximately 70 with no averaging. This is lower than for the oil experiment due to smaller sample volume (5 mm tube instead of 10 mm tube). Figure 15 compares data collected on

Probe	τ_{min}	τ_{max}	τ steps	$\pi/2$ pulse	π pulse	Pulse length	t_{echo}	Echoes	Averages
10 mm	1 ms	11000 ms	15	1 W	3.98 W	23 μ s	400 μ s	2000	1

Table 3: Pulse sequence parameters for the cooking oil experiments.

Probe	τ_{min}	τ_{max}	τ steps	$\pi/2$ pulse	π pulse	Pulse length	t_{echo}	Echoes	Averages
5 mm	100 ms	700 ms	20	1.58 W	6.38 W	38 μ s	400 μ s	2000	1

Table 4: Pulse sequence parameters for the milk experiments.

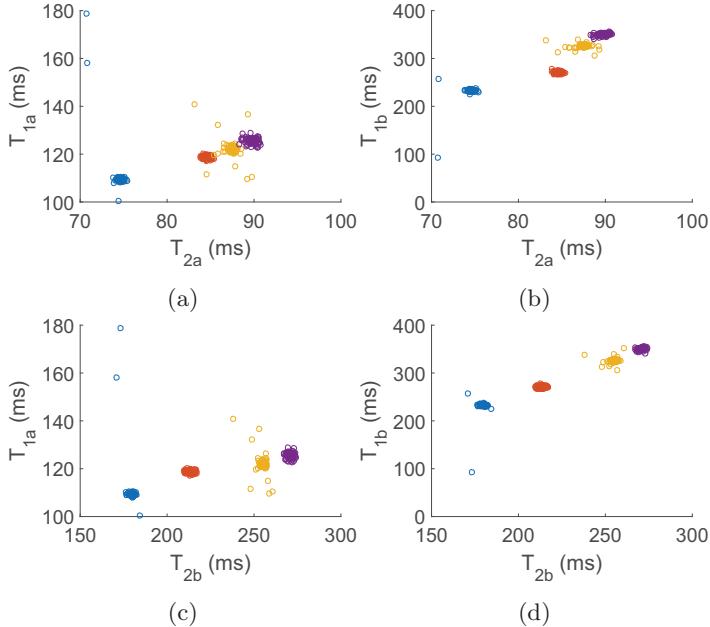


Figure 11: Comparison of different features extracted for classification. The blue, orange, yellow, and purple circles correspond to olive, canola, corn, and vegetable oil, respectively. The y -axes of (a) and (b) show the first and second components of T_1 , respectively; the x -axes of both plots show the first component of T_2 . The y -axes of (c) and (d) are the same as for (a) and (b), but the x -axis now shows the second component of T_2 .

fresh milk at 9 AM on the first day with data collected on the same sample at 9 AM on the third day. The plot shows that the distribution spreads along the T_1 dimension as the milk spoils, while its peak value remains almost constant. Thus, a fraction of the protons develop a larger T_1/T_2 ratio as the milk spoils, suggesting that their motion is slowing down. This is probably caused by coagulation of the milk proteins, which is in agreement with visual inspection of the sample.

In order to extract information about changes in the T_1 distribution, we treat it as a bi-Gaussian function and observe changes in the peak values versus time. The T_2 distribution does not change significantly over time, and is well-fitted by a Gaussian. The fitted values reveal a significant

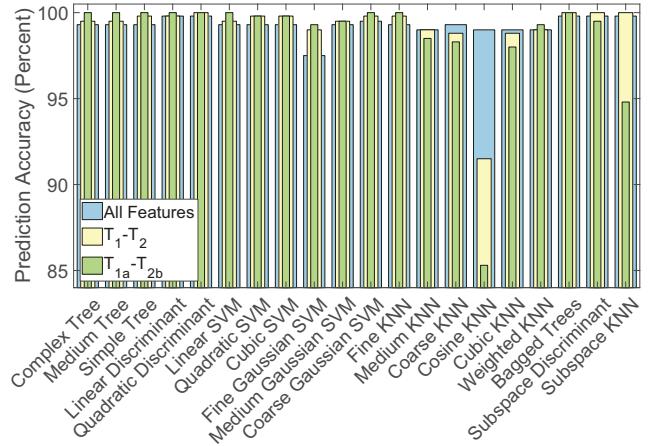


Figure 12: Prediction accuracy of all classification algorithms on the cooking oil samples. Results are shown for models trained on three different combinations of classification features.

shift in the first T_1 component between the first two data sets and the rest, as seen in Figure 16. Thus, this change is likely to be correlated with spoilage; the dashed lines clearly separate the fresh (left) and spoiled (right) samples as determined by visual inspection. There is also some uncorrelated variation of T_2 and the second T_1 component, which may be due to small shifts in sample temperature.

For classification, we use only the T_1a parameter for training the model as the other parameters do not seem to be correlated with spoilage. We manually label the first two data sets as fresh and the rest as spoiled, and then train the same models used in the oil experiments. The accuracy of a subset of the models was evaluated using cross validation. Figures 18, 17, and 19 summarize the prediction accuracy, training time, and prediction speed, respectively. We show a full summary of the results in Appendix B. As expected, the algorithms shown have a prediction accuracy of 100%. The simple tree algorithm is the most efficient in this case with 100% accuracy, 2400 obs/s prediction speed, and 0.27 sec training time.

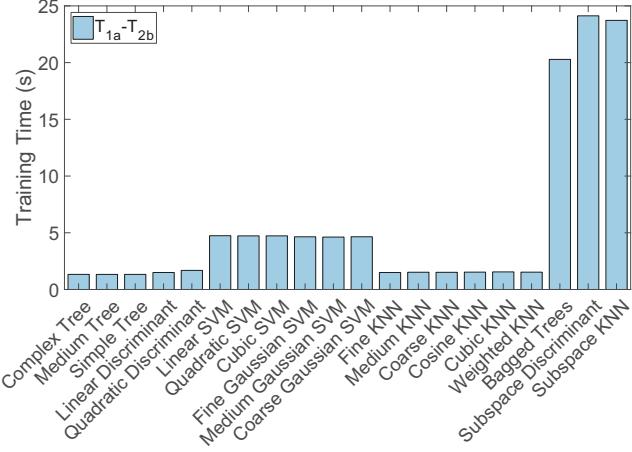


Figure 13: Training time of all classification algorithms tested on the cooking oil samples. Results are shown for models trained on the T_{1a} - T_{2b} feature set.

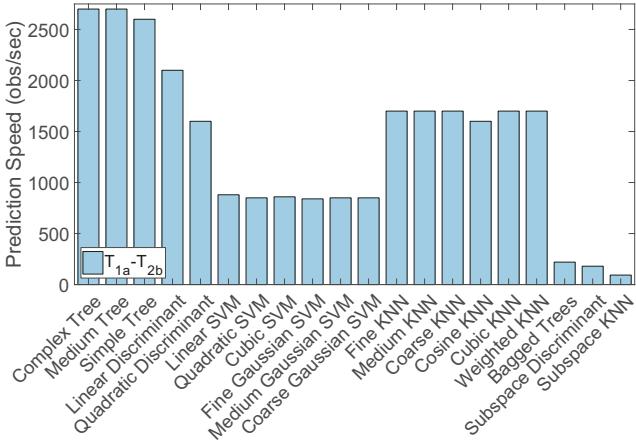


Figure 14: Prediction speed in obs/s of each algorithm tested on the cooking oil samples. Results are shown for models trained on the T_{1a} - T_{2b} feature set.

3.6. Sodium Experiments

We investigated the feasibility of improving classification accuracy by measuring both the proton and sodium T_1 - T_2 maps of several soy sauce samples (regular, less sodium, ponzu, and teriyaki). In each case, 20 ^{23}Na and 5 ^1H data sets of data were collected. The ambient temperature was maintained at 30 °C (303 K) to ensure valid comparisons between the data sets.

We initially focus on the sodium data set. As in the previous experiments, we extracted five features from the data for classification. However, we apply a different approach to find the values of these features. The correlation maps are highly localized, so T_1 and T_2 can be considered to be generated by a single component. Therefore we can fit the maps using two-dimensional non-singular normal distributions. In addition, we can predict the mean values of T_1 and T_2 by independently estimating their one-dimensional

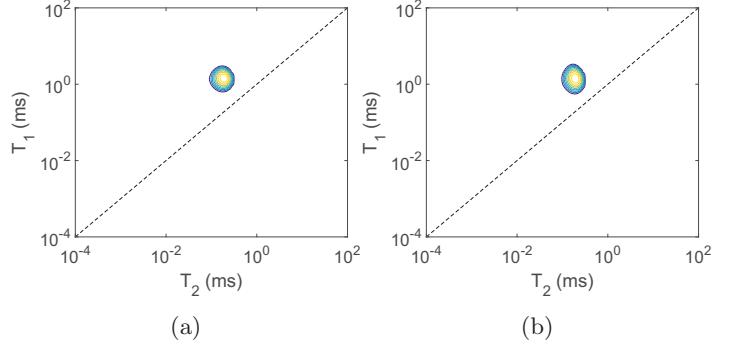


Figure 15: A comparison of the T_1 - T_2 maps collected on the same milk sample at (a) 9 AM on the first day and (b) 9 AM on the third day. The T_1 distribution spreads as the milk spoils. Each data set has an SNR of approximately 70 with no averaging.

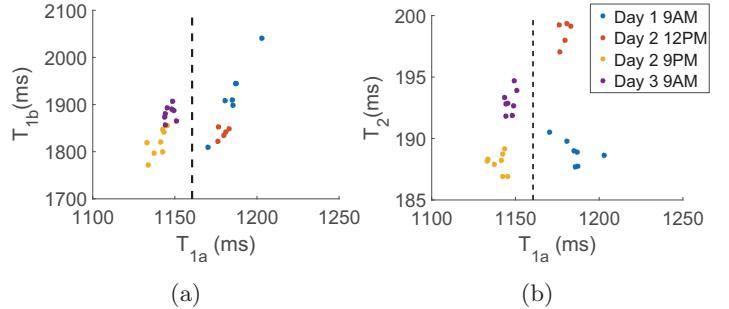


Figure 16: Classification features extracted from the milk sample; the dashed lines separate spoiled milk (left) and fresh milk (right). (a) The first T_1 component versus the second T_1 component. (b) The T_2 component versus the second T_1 component.

distributions to improve model accuracy.

The collected data is summarized in Figure 20 and a table of the pulse sequence parameters can be found in Tables 5 and 6. The samples tends to form tight clusters and there is distinct separation between each of them; ponzu sauce is particularly distinct. Two of the extracted parameters, the mean values of T_1 and T_2 , are shown in Table 7.

Classification using the selected features was carried out as in the earlier experiments, and 50-fold cross validation was again used to estimate prediction accuracy. The training time, prediction accuracy, and prediction speed of various algorithms trained on three different sets of features are compared in Figures 21, 22, and 23. A full summary of the results is shown in Appendix B. The plots show that almost all the information is contained in the mean T_1 and T_2 values, as for the earlier samples. This is probably because the covariance parameters are highly dependent on the amount of regularization and are thus not sample-specific. Thus, for some algorithms, accuracy goes up significantly if only the mean T_1 and T_2 are used as features. Overall, the simple tree is the most efficient algorithm for this data set.

The basic setup of the sodium D - T_2 experiment is as same as for the T_1 - T_2 experiment. The diffusion encoding param-

Probe	τ_{min}	τ_{max}	τ steps	t_{echo}	$\pi/2$ pulse	π pulse	Pulse length	Echoes	Averages
5 mm	1 ms	1000 ms	20	300 μ s	3.16 W	12.59 W	35 μ s	150	128

Table 5: Pulse sequence parameters for the T_1 - T_2 sodium experiments.

Probe	δ	Δ	Gradient steps	t_{echo}	$\pi/2$ pulse	π pulse	Pulse length	Echoes	Averages
5 mm	7000 μ s	9000 μ s	20	300 μ s	6.31 W	25.12 W	17 μ s	250	1024

Table 6: Pulse sequence parameters for the Diffusion- T_2 sodium experiments.

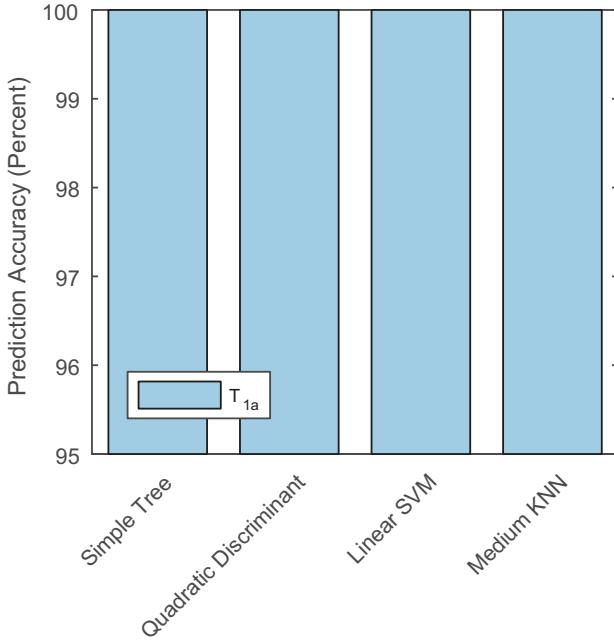


Figure 17: Prediction accuracy on a subset of classification algorithms tested on the milk sample. Results are shown for models trained on only the T_{1a} feature.

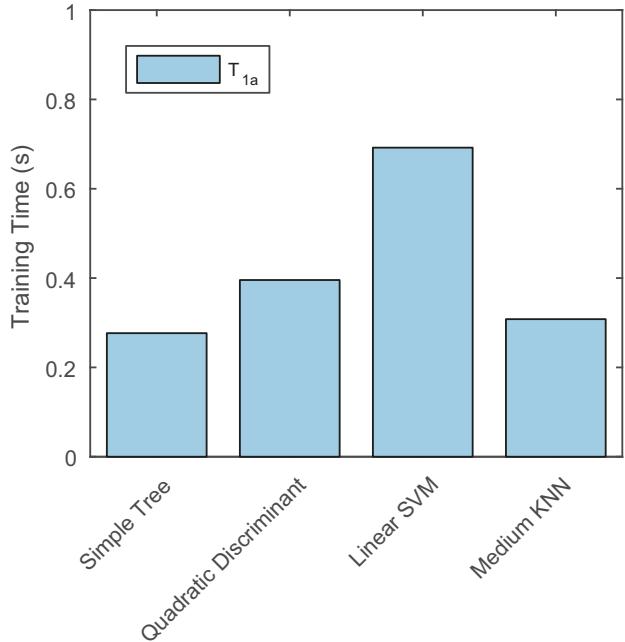


Figure 18: Training time on a subset of classification algorithms tested on the milk sample. Results are shown for models trained on only the T_{1a} feature.

eters in the pulse sequence (Δ and δ) are modified according to the tested sample to optimize the SNR. We observe that all the samples tend to converge into tight clusters, as shown in Figure 24a. The mean T_2 values of the samples match those obtained from T_1 - T_2 experiments. The mean diffusion constant D is somewhat more widely scattered, mainly because of limited SNR due to hardware constraints (e.g., sample volume, maximum available gradient strength). The problem can be resolved in several ways such as controlling the ambient temperature, using post processing techniques to compensate for temperature shifts, and improving SNR per scan to accelerate the experiments. Nevertheless, D is still able to assist in distinguishing between the different sauce samples, and is thus marked as an important feature for classification.

The proton T_1 - T_2 map for the same soy sauce samples is

shown in Figure 24b. The plot shows that ^1H T_1 and T_2 change even more dramatically between the different samples than for ^{23}Na . This may be because of the absence of quadrupolar relaxation for protons. Another possibility is that the protons are bound to molecular aggregates that span a large range of sizes (and thus rotational correlation times) between the samples, while the sodium is predominantly in solution (in ionized form) for all samples. However, the overall trend in sample properties remains the same: regular soy sauce has the lowest T_1 and T_2 values, while ponzu sauce has the highest values. Moreover, in both cases ponzu is widely separated from the other samples. We conclude that ^1H T_1 and T_2 can also provide useful classification features for this data set.

Total sodium content was also estimated via a one-point calibration with a reference brine solution. The estimated

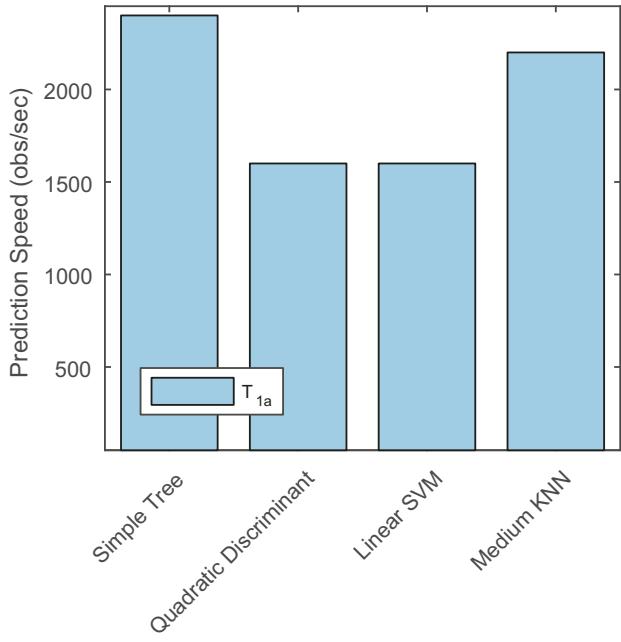


Figure 19: Prediction speed on a subset of classification algorithms tested on the milk sample. Results are shown for models trained on only the T_{1a} feature.

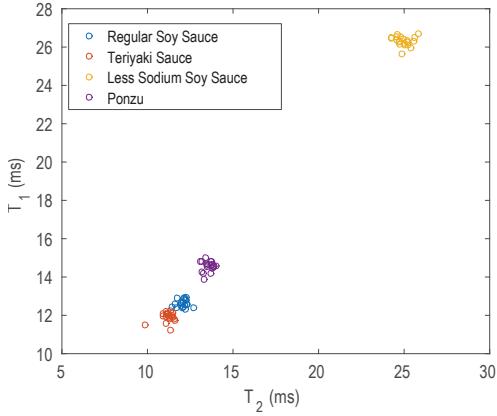


Figure 20: ^{23}Na T_1 - T_2 correlation maps collected from four different kinds of soy sauce. SNR ≈ 35 for each set of experiments after 128 scans.

values (per teaspoon) for less sodium soy sauce, regular soy sauce, ponzu sauce, and teriyaki sauce are 196, 305, 158, and 204 mg, respectively. These values are in excellent agreement with those on the product labels, which are 192, 307, 177, and 203 mg, respectively.

In conclusion, we can clearly classify the various sauce samples using sodium and proton NMR measurements: they can be effectively classified using either sodium T_1 and T_2 , sodium D , or proton T_1 and T_2 . Combinations of these features can also be measured to improve classification accuracy if needed, at the cost of longer data collection and

Soy Sauce	T_1 (ms)	T_2 (ms)
Regular	11.91	11.20
Less Sodium	14.56	13.57
Teriyaki	12.62	12.07
Ponzu	26.31	24.96

Table 7: Mean ^{23}Na T_1 and T_2 values of the soy sauce samples extracted from the T_1 - T_2 correlation maps.

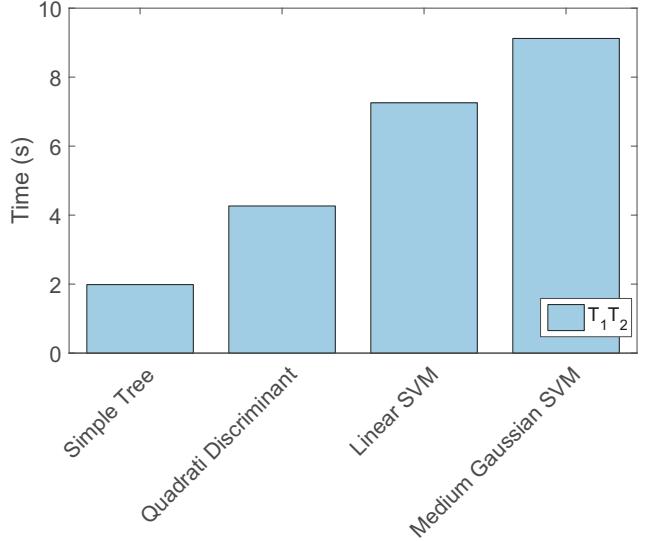


Figure 21: Training time on a subset of classification algorithms tested on the soy sauce samples. Results are shown for models trained on only the T_1T_2 feature.

model training times. The experiments thus demonstrate the possibilities of authenticating unknown sauce samples and characterizing their freshness.

3.7. Temperature-controlled experiments

Now we consider temperature effects upon the experiments. If there is no temperature control mechanism, the field will drift with time as the room temperature changes. The high temperature coefficient (about $-1.1\text{%/}^\circ\text{C}$) of the NdFeB permanent magnet will then result in significant changes in the Larmor frequency over time. For example, the proton frequency decreases by 24 kHz/ $^\circ\text{C}$. Such drifts must be taken into account while running long experiments.

We used regular soy sauce as an example to demonstrate the effects of temperature upon our measurement results. In order to exclude the influence of poorly-controlled room temperature fluctuations, both sample and magnet were placed inside a Peltier-controlled incubator (see Figure 6). The sample and magnet temperature was then varied from 22°C to 34°C (295 K to 307 K) in 2°C steps, and T_1 - T_2 correlation plots were generated at each step. The collected data is summarized in Figure 25. The data shows that both T_1

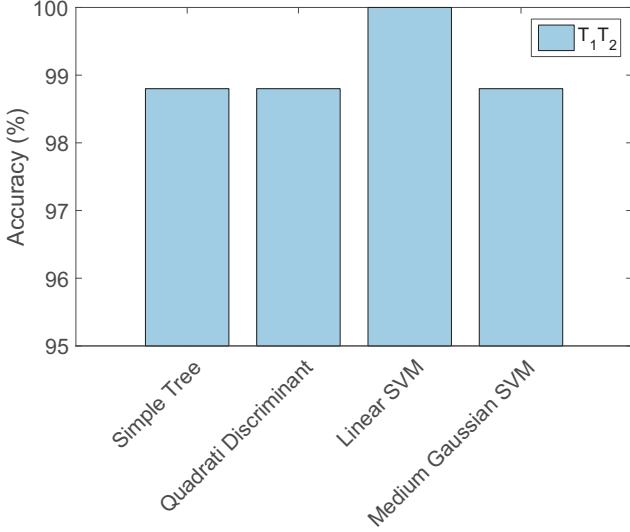


Figure 22: Prediction accuracy on a subset of classification algorithms tested on the soy sauce samples. Results are shown for models trained on only the T_1T_2 feature.

and T_2 are strongly dependent on temperature. Moreover, they increase linearly with temperature within this limited range. Thus, we expect temperature drifts to have a strong influence on our classification results.

Indeed, the accuracy of predicting an unknown sample drops significantly if the temperature-dependent data set is fed into the trained model generated during the T_1 - T_2 experiment, as summarized in Figure 26 (a complete set of results from all 20 classification methods is shown in Appendix B). Taking the linear SVM classification method as an example, the accuracy drops from 100.0% to 71.4%. Notice that the mean values of T_1 and T_2 for regular soy sauce are the lowest among four samples. This simplifies the classification difficulty and improves the result because the data is unlikely to be mixed up with data from other samples. In other words, the classification results would have been even worse if temperature-dependent data from the other samples had been fed into the same trained model. Thus, we have to take sample temperature into account during classification, and the best results will be obtained if the temperature is kept constant.

4. Summary and conclusions

We have explored the potential of classifying different kinds of food products using a double-tuned, low-field NMR system. The system can reliably distinguish the tested samples, including cooking oils, milk, and soy sauces. The design of the custom gradient coil used by the system has been analyzed in detail. The system uses an easy-to-use and efficient double-tuned impedance matching network for easily interleaving proton and sodium NMR experiments. The numerical optimization of this network was presented in detail;

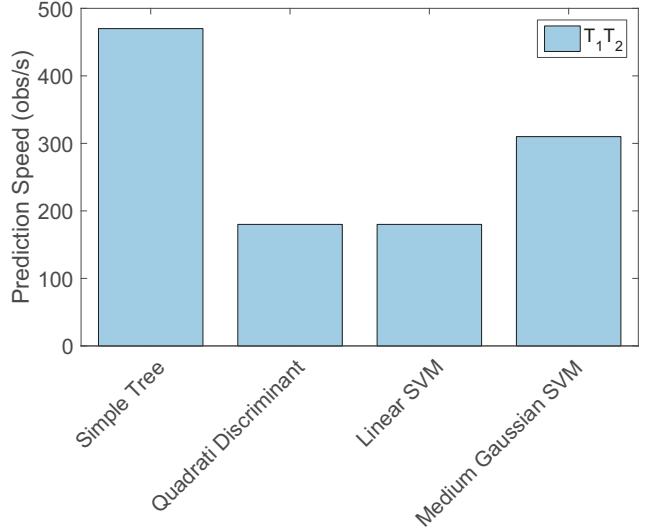


Figure 23: Prediction speed on a subset of classification algorithms tested on the soy sauce samples. Results are shown for models trained on only the T_1T_2 feature.

the results demonstrate that the RF coil can be effectively double-tuned with an acceptable amount of power loss.

The system was used to measure T_1 - T_2 and D - T_2 correlation data for both ^1H and ^{23}Na . Important features were extracted from these data sets and used to train a variety of classification models. Many of the trained models were shown to have a classification accuracy of as high as 100% for different food products. We also showed that sample temperature has a major effect on classification accuracy, which highlights the importance of temperature control for obtaining reliable results. In sum, the results strongly support the capability of ^1H and ^{23}Na 2D low-field NMR to be used in quality assessment of food products. The technique is promising for developing robust prediction models and databases of authentic products. In addition, integration of NMR with other analytical measurement techniques such as NIR spectroscopy are of interest for future work.

Acknowledgements

The authors would like to acknowledge the National Science Foundation (NSF) for funding support. This work was partially funded by the NSF through grant CCF-1563688.

Appendix A. Derivation of Eqn. (2)

We start by writing down the Biot-Savart law for a single wire:

$$\vec{B}(\vec{r}) = \frac{\mu_0 I}{4\pi} \int \frac{d\vec{s} \times \vec{r}}{r^3}, \quad (\text{A.1})$$

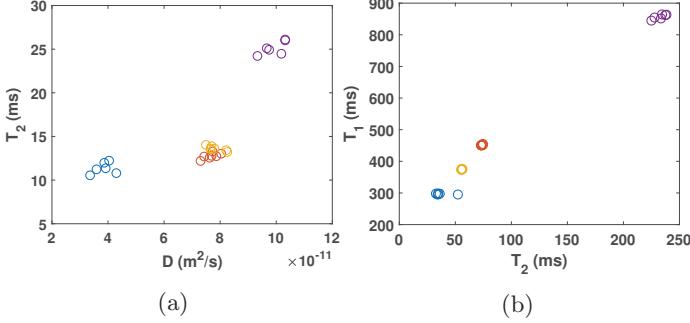


Figure 24: (a) Soy sauce ^{23}Na D - T_2 map. The blue, orange, yellow, and purple circles correspond to regular soy, less sodium soy, teriyaki, and ponzu sauce, respectively. The mean values of T_2 matches results from the T_1 - T_2 map; the diffusion constants of different samples differ significantly from each other. SNR ≈ 15 for each set of experiments after 1024 scans. (b) Soy sauce ^1H T_1 - T_2 map. Different samples are easily distinguished in this case as well. SNR ≈ 125 for each set of experiments after 8 scans.

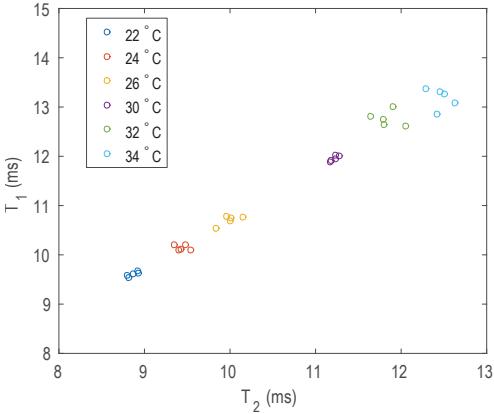


Figure 25: Temperature shift experiment for the regular soy sauce sample. Five sets of data were collected at each temperature. Over this limited range, both T_1 and T_2 increase linearly with temperature.

where $d\vec{s}$ is the differential vector along the wire, \vec{r} is the radius vector from the wire to the point at which the field is calculated, and I is the current. If the wire of interest crosses the y - z plane at the point (c, d) and we assume the magnetic field produced by the wire is much smaller than the static field along the z -direction, then only the z -component of $\vec{B}(\vec{r})$ is of interest. We can find this quantity by integrating (A.1) along the x -axis:

$$B_z(y, z) = \frac{\mu_0 I}{2\pi} \frac{(c - y)}{(c - y)^2 + (d - z)^2} = \frac{\mu_0 I}{2\pi} \Re \left[\frac{1}{ae^{i\phi} - \xi} \right], \quad (\text{A.2})$$

where the complex variables $ae^{i\phi} \equiv (c + id)$ and $\xi \equiv (y + iz)$ are defined for convenience. In a small region around the origin (near the RF coil), $|\xi| \ll a$. Because of this fact, we

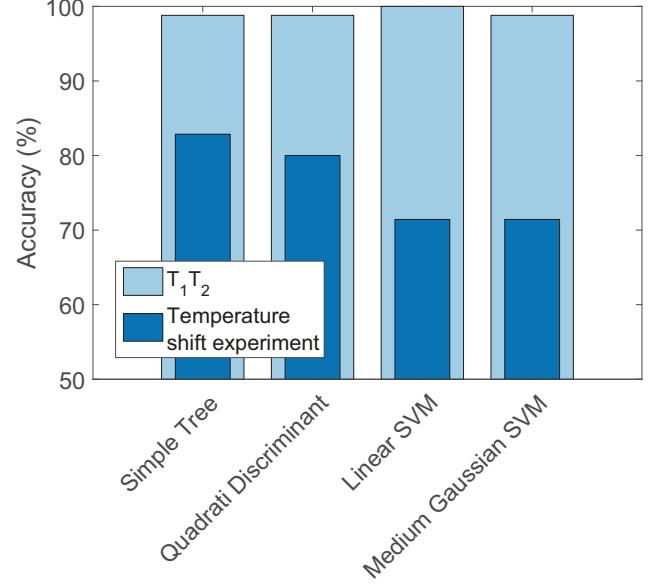


Figure 26: Comparison of prediction accuracy for the regular soy sauce sample i) the original T_1 - T_2 map data (measured at constant temperature), and ii) the temperature-dependent data. Prediction accuracy drops significantly in the latter case.

can expand the denominator as a power series:

$$B_z(\xi) = \frac{\mu_0 I}{2\pi a} \Re \left[\frac{e^{-i\phi}}{1 - \frac{\xi}{a} e^{-i\phi}} \right] \approx \frac{\mu_0 I}{2\pi a} \Re \left[\sum_{n=0}^{\infty} \frac{\xi^n}{a^n} e^{-i(n+1)\phi} \right]. \quad (\text{A.3})$$

To calculate the total field, we add up the contribution from each wire using the superposition principle. The four wires are located at $\phi, \phi + \frac{\pi}{2}, \phi + \pi$, and $\phi + \frac{3\pi}{2}$. This results in the terms corresponding to $n = [1, 5, 9, \dots]$ adding together and all other terms canceling, giving us a total field described by

$$B_z(\xi) \approx \frac{2\mu_0 I}{\pi a} \Re \left[e^{-2i\phi} \left(\frac{\xi}{a} \right) + e^{-6i\phi} \left(\frac{\xi}{a} \right)^5 + e^{-10i\phi} \left(\frac{\xi}{a} \right)^9 + \dots \right]. \quad (\text{A.4})$$

At the RF coil near the origin, the first term will dominate, resulting in a magnetic field of

$$B_z(y, z) \approx \frac{2\mu_0 I}{\pi a^2} [\cos(2\phi)y + \sin(2\phi)z]. \quad (\text{A.5})$$

Taking the derivative of (A.5), we get the y and z gradients shown below:

$$g_y \equiv \frac{\delta B_z}{\delta y} = \left(\frac{2\mu_0 I}{\pi a^2} \right) \cos(2\phi) \quad (\text{A.6})$$

$$g_z \equiv \frac{\delta B_z}{\delta z} = \left(\frac{2\mu_0 I}{\pi a^2} \right) \sin(2\phi).$$

We can solve for the magnetic field of the distorted case using the same techniques that were used for the ideal geometry due to the fact that the wires still lie on the circumference of a circle. However, the wires now have arbitrary vertical and horizontal separations of $2a_1$ and $2a_2$. Thus, the coordinates of the wires on the (y, z) plane are $(a_1, a_2), (a_1, -a_2), (-a_1, a_2)$, and $(-a_1, -a_2)$. These coordinates lie on a circle of radius $a = \sqrt{a_1^2 + a_2^2}$ and each coordinate has an azimuthal angle of $\phi, -\phi, \phi + \pi$, and $-\phi + \pi$, respectively, where $\phi \equiv \tan^{-1}(a_2/a_1)$.

Applying the same techniques, we start by solving for the field near the origin:

$$\begin{aligned}
B_z(\xi) &\approx \frac{\mu_0 I}{\pi a} \Re \left[(e^{-2i\phi} - e^{2i\phi}) \left(\frac{\xi}{a} \right) \right. \\
&\quad \left. + (e^{-4i\phi} - e^{4i\phi}) \left(\frac{\xi}{a} \right)^3 + (e^{-6i\phi} - e^{6i\phi}) \left(\frac{\xi}{a} \right)^5 + \dots \right] \\
&= \frac{\mu_0 I}{\pi a} \Re \left[-2i \sin(2\phi) \left(\frac{\xi}{a} \right) - 2i \sin(4\phi) \left(\frac{\xi}{a} \right)^3 \right. \\
&\quad \left. - 2i \sin(6\phi) \left(\frac{\xi}{a} \right)^5 + \dots \right] \\
&= \frac{2\mu_0 I}{\pi a^2} \left[\sin(2\phi) z + \sin(4\phi) \frac{(3y^2 z - z^3)}{a^2} + \dots \right]. \tag{A.7}
\end{aligned}$$

From eqn. (A.7), it can be seen that fewer terms cancel in the power series expansion. Specifically, in the distorted geometry, every other term will survive while in the ideal one, only every fourth term survives. This is because the distorted geometry has less symmetry than the ideal geometry.

We can now take the partial derivatives of eqn. (A.7) to get the following expression for the gradients:

$$\begin{aligned}
g_y &\equiv \frac{\delta B_z}{\delta y} \approx \left(\frac{2\mu_0 I}{\pi a^2} \right) \left[6 \sin(4\phi) \frac{yz}{a^2} + \dots \right], \tag{A.8} \\
g_z &\equiv \frac{\delta B_z}{\delta z} \approx \left(\frac{2\mu_0 I}{\pi a^2} \right) \left[\sin(2\phi) + 3 \sin(4\phi) \frac{(y^2 - z^2)}{a^2} + \dots \right].
\end{aligned}$$

Appendix B. Complete Summary of Classification Data

For completeness, this appendix summarizes classification results obtained for various samples using all 20 available algorithms. Figs. B.27, B.28 and B.29 respectively summarize prediction accuracy, training time, and prediction speed for the milk sample.

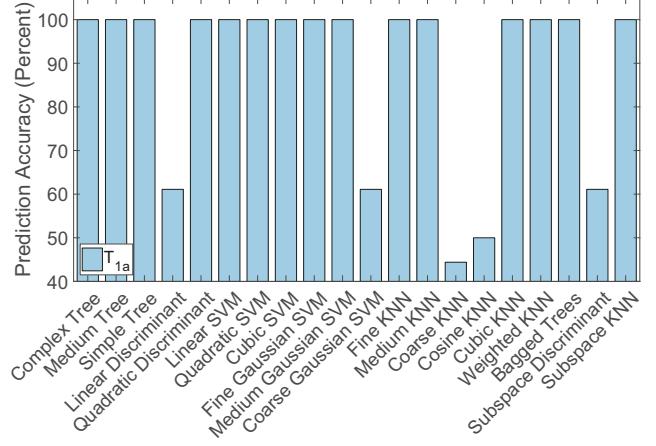


Figure B.27: Prediction accuracy of all classification algorithms tested on the milk sample. Results are shown for models trained on only the T_{1a} feature.

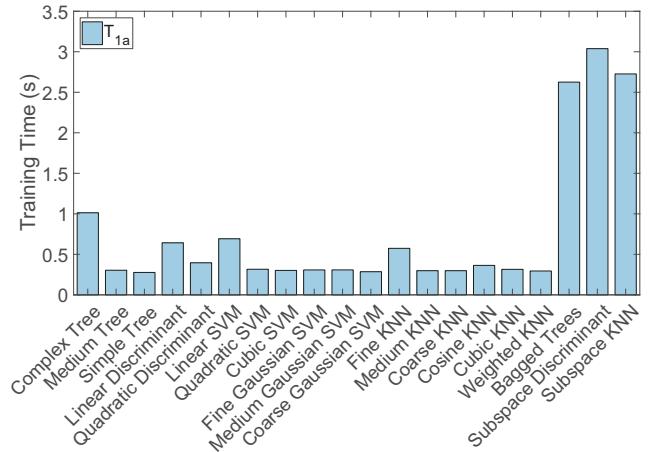


Figure B.28: Training time of all classification algorithms tested on the milk sample. Results are shown for models trained on only the T_{1a} feature.

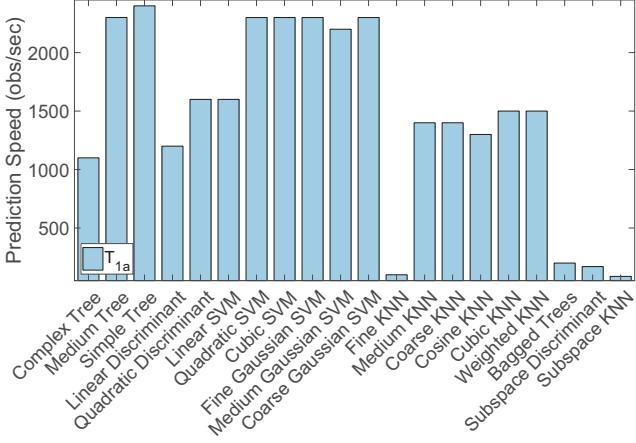


Figure B.29: Prediction speed of all classification algorithms tested on the milk sample. Results are shown for models trained on only the T_{1a} feature.

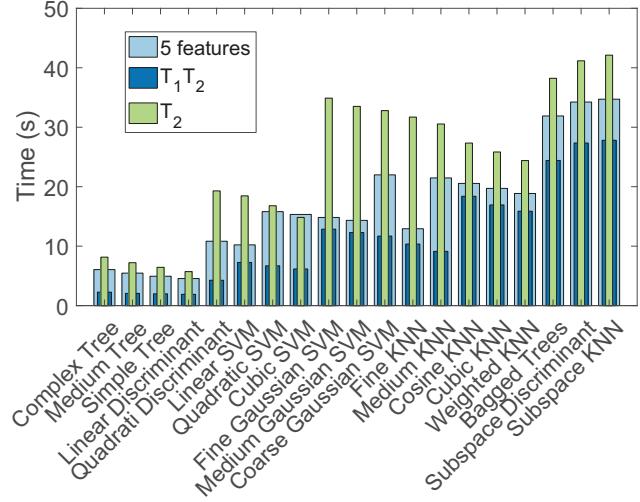


Figure B.31: Training time of all classification algorithms tested on the soy sauce samples. Results are shown for models trained on three combinations of features.

Figs. B.30, B.31 and B.32 respectively summarize prediction accuracy, training time, and prediction speed for the soy sauce samples.

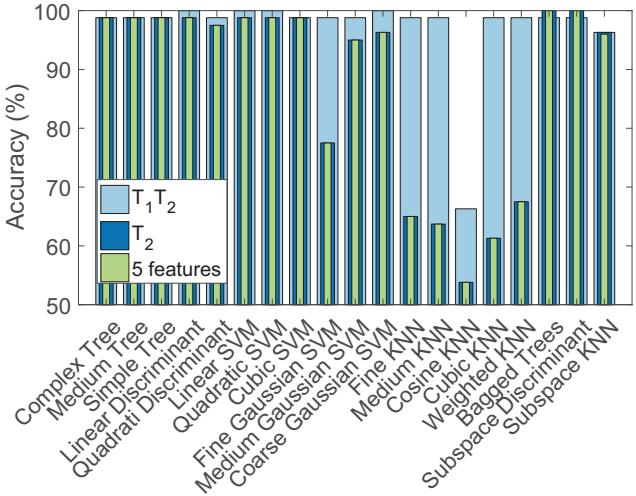


Figure B.30: Prediction accuracy of all classification algorithms tested on the soy sauce samples. Results are shown for models trained on three combinations of features.

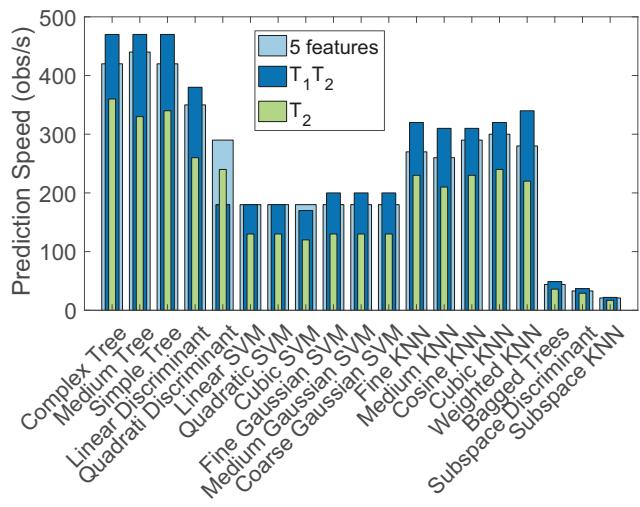


Figure B.32: Prediction speed of all classification algorithms tested on the soy sauce samples. Results are shown for models trained on three combinations of features.

Fig. B.33 summarizes prediction accuracy for the regular soy sauce sample using the original and temperature-dependent data sets.

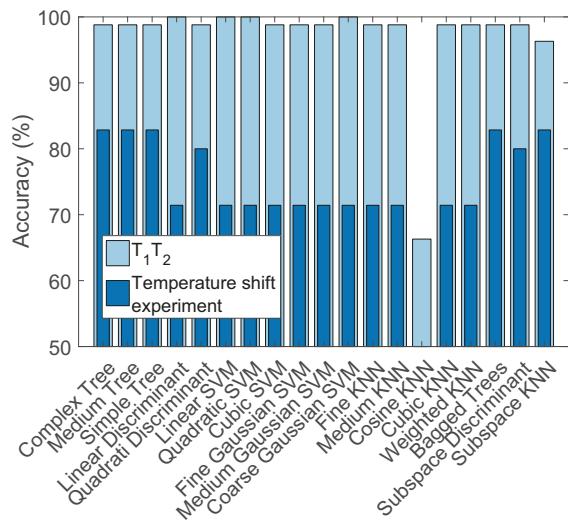


Figure B.33: Comparison of prediction accuracy for the regular soy sauce sample using i) the original T_1-T_2 map data (measured at constant temperature), and ii) the temperature-dependent data.

References

[1] David I. Ellis, Victoria L. Brewster, Warwick B. Dunn, J. William Allwood, Alexander P. Golovanov, and Royston Goodacre. Fingerprinting food: current technologies for the detection of food adulteration and contamination. *Chemical Society Reviews*, 41(17):5706–5727, 2012.

[2] Albert I. Wertheimer and Jeremiah Norris. Safeguarding against substandard/counterfeit drugs: mitigating a macroeconomic pandemic. *Research in Social and Administrative Pharmacy*, 5(1):4–16, 2009.

[3] Cheng Chen, Fengchao Zhang, Swarup Bhunia, and Soumyajit Mandal. Broadband quantitative nqr for authentication of vitamins and dietary supplements. *Journal of Magnetic Resonance*, 278:67 – 79, 2017.

[4] F. Zhang, N. V. R. Masna, S. Bhunia, C. Chen, and S. Mandal. Authentication and traceability of food products through the supply chain using nqr spectroscopy. In *2017 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pages 1–4, Oct 2017.

[5] Zheng Xu, Robert H. Morris, Martin Bencsik, and Michael I. Newton. Detection of virgin olive oil adulteration using low field unilateral NMR. *Sensors (Basel)*, 14(2):2028–2035, Feb 2014. sensors-14-02028[PII].

[6] Cleiton A. Nunes. Vibrational spectroscopy and chemometrics to assess authenticity, adulteration and intrinsic quality parameters of edible oils and fats. *Food Research International*, 60:255 – 261, 2014. Authenticity, Typicality, Traceability and Intrinsic Quality of Food Products.

[7] Raffaele Sacchi, Francesco Addeo, and Livio Paolillo. ^1H and ^{13}C NMR of virgin olive oil. an overview. *Magnetic Resonance in Chemistry*, 35(13):S133–S145, 1997.

[8] S. Ok. Detection of olive oil adulteration by low-field NMR relaxometry and UV-Vis spectroscopy upon mixing olive oil with various edible oils. *Grasas y Aceites*, 68(1):173, 2017.

[9] B. S. Luh. Industrial production of soy sauce. *Journal of Industrial Microbiology*, 14(6):467–471, Jun 1995.

[10] Tian Luo, Yunfei Xie, Yue Dong, Aihui Liu, and Yuming Dong. Quality assessment of soy sauce using underivatized amino acids by capillary electrophoresis. *International Journal of Food Properties*, 0(0):1–10, 2018.

[11] T. Natcha and S. Panmanas. Detection of adulteration of soy sauce by brine using near infrared spectroscopy. *Agriculture and Agricultural Science Procedia*, 2:212 – 217, 2014. 2nd International Conference on Agricultural and Food Engineering (CAFEi 2014) - New Trends Forward.

[12] Yonghong Zhu, Yan Yang, Zhaoxu Zhou, Genrong Li, Mei Jiang, Chun Zhang, and Shiqi Chen. Direct determination of free tryptophan contents in soy sauces and its application as an index of soy sauce adulteration. *Food Chemistry*, 118(1):159 – 162, 2010.

[13] Bong-Kuk Ko, Hyuk-Jin Ahn, Frans van den Berg, Cherl-Ho Lee, and Young-Shick Hong. Metabolomic insight into soy sauce through ^1H nmr spectroscopy. *Journal of Agricultural and Food Chemistry*, 57(15):6862–6870, 2009. PMID: 19591484.

[14] Ghulam Mustafa Kamal, Xiaohua Wang, Bin Yuan, Jie Wang, Peng Sun, Xu Zhang, and Maili Liu. Compositional differences among chinese soy sauce types studied by ^{13}C NMR spectroscopy coupled with multivariate statistical analysis. *Talanta*, 158:89 – 99, 2016.

[15] Michael Lu, Yvonne Shiao, Jacklyn Wong, Raishay Lin, Hannah Kravis, Thomas Blackmon, Tanya Pakzad, Tiffany Jen, Amy Cheng, Jonathan Chang, et al. Milk spoilage: Methods and practices of detecting milk quality. *Food and Nutrition Sciences*, 4(07):113, 2013.

[16] Hilde M stlie, Merete H Helland, and Judith A Narvhus. Growth and metabolism of selected strains of probiotic bacteria in milk. *International Journal of Food Microbiology*, 87(1):17 – 27, 2003.

[17] Yuan-Geuy Lee, Hung-Yi Wu, Chuan-Liang Hsu, Hong-Jen Liang, Chiun-Jye Yuan, and Hung-Der Jang. A rapid and selective method for monitoring the growth of coliforms in milk using the combination of amperometric sensor and reducing of methylene blue. *Sensors and Actuators B: Chemical*, 141(2):575–580, 2009.

[18] Sijing Huang, Shutian Ge, Lingwen He, Qingyun Cai, and Craig A Grimes. A remote-query sensor for predictive indication of milk spoilage. *Biosensors & bioelectronics*, 23(11):17451748, June 2008.

[19] John Erik Haugen, Knut Rudi, Solveig Langsrud, and Sylvia Bredholt. Application of gas-sensor array technology for detection and monitoring of growth of spoilage bacteria in milk: A model study. *Analytica Chimica Acta*, 565(1):10–16, 2006.

[20] H. M. Al-Qadiri, M. Lin, M. A. Al-Holy, A. G. Cavinato, and Barbara A. Rasco. Monitoring quality loss of pasteurized skim milk using visible and short wavelength near-infrared spectroscopy and multivariate analysis. *Journal of dairy science*, 91(3):950–958, 2008.

[21] Sanaa O. Yagoub, Fardous A. Bellow, and I. El Zubeir. Effect of temperature and storage period on the constituents of milk inoculated with *pseudomonas aeruginosa*. *Research Journal of Microbiology*, 3(1):30–34, 2008.

[22] Yu. B. Monakhova, T. Kuballa, and D. W. Lachenmeier. Chemometric methods in nmr spectroscopic analysis of food products. *Journal of Analytical Chemistry*, 68(9):755–766, Sep 2013.

[23] Y.-Q. Song, L. Venkataraman, M. D. Hürlimann, M. Flbaum, P. Frulla, and C. Straley. T1–T2 correlation spectra obtained using a fast two-dimensional Laplace inversion. *Journal of Magnetic Resonance*, 154(2):261–268, 2002.

[24] D. Ariando, D. Mirando, J. Shi, M. Greer, C. Chen, and S. Mandal. A low-cost miniature nmr spectrometer for active learning. In *Experimental NMR Conference (ENC)*, Mar 2017.

[25] I. Zupancic and J. Pirs. Coils producing a magnetic field gradient for diffusion measurements with NMR. *Journal of Physics E: Scientific Instruments*, 9(1):79, 1976.

[26] M. D. Schnall, V. Harihara Subramanian, J. S. Leigh Jr., and B. Chance. A new double-tuned probe for concurrent ^1H and ^{31}P nmr. *Journal of Magnetic Resonance (1969)*, 65(1):122–129, 1985.

[27] Dongjiang Qiao, Robert Molfino, Steven M Lardizabal, Brandon Pillans, Peter M Asbeck, and George Jerinic. An intelligently controlled RF power amplifier with a reconfigurable MEMS-varactor tuner. *IEEE Transactions on Microwave Theory and Techniques*, 53(3):1089–1095, 2005.

[28] J. P. Butler, J. A. Reeds, and S. V. Dawson. Estimating solutions of first kind integral equations with nonnegative constraints and optimal smoothing. *SIAM Journal on Numerical Analysis*, 18(3):381–397, 1981.

[29] Georgios P. Danezis, Aristidis S. Tsagkaris, Vladimir Brusic, and Constantinos A. Georgiou. Food authentication: state of the art and prospects. *Current Opinion in Food Science*, 10:22–31, 2016.

[30] Camila Maione and Rommel Melgaço Barbosa. Recent applications of multivariate data analysis methods in the authentication of rice and the most analyzed parameters: A review. *Critical reviews in food science and nutrition*, pages 1–12, 2018.

[31] Dong Kyu Lim, Nguyen Phuoc Long, Changyeun Mo, Ziyuan Dong, Lingmei Cui, Giyoung Kim, and Sung Won Kwon. Combination of mass spectrometry-based targeted lipidomics and supervised machine learning algorithms in detecting adulterated admixtures of white rice. *Food Research International*, 100:814–821, 2017.

[32] Piotr S. Gromski, Howbeer Muhamadali, David I. Ellis, Yun Xu, Elon Correa, Michael L. Turner, and Royston Goodacre. A tutorial review: Metabolomics and partial least squares-discriminant analysis a marriage of convenience or a shotgun wedding. *Analytica Chimica Acta*, 879:10 – 23, 2015.

[33] Ji-Hyun Kim. Estimating classification error rate: Repeated cross-validation, repeated hold-out and bootstrap. *Computational Statistics & Data Analysis*, 53(11):3735 – 3745, 2009.

[34] Sotiris B Kotsiantis, I Zaharakis, and P Pintelas. Supervised machine learning: A review of classification techniques. *Emerging artificial intelligence applications in computer engineering*, 160:3–24, 2007.

[35] S. Mandal and Y.-Q. Song. Heteronuclear J-coupling measurements in grossly inhomogeneous magnetic fields. *Journal of Magnetic Resonance*, 255:15–27, 2015.