Direct Imaging of Functional Networks

Eric C. Wong

Abstract

In blood-oxygenation-level-dependent functional magnetic resonance imaging (fMRI), current methods typically acquire $\sim 500,000$ imaging voxels at each time point, and then use computer algorithms to reduce this data to the coefficients of a few hundred parcels or networks. This suggests that the amount of relevant information present in the fMRI signal is relatively small, and presents an opportunity to greatly improve the speed and signal to noise ratio (SNR) of the fMRI process. In this work, a theoretical framework is presented for calculating the coefficients of functional networks directly from highly undersampled fMRI data. Using predefined functional parcellations or networks and a compact k-space trajectory that samples data at optimal spatial scales, the problem of estimating network coefficients is reformulated to allow for direct least squares estimation, without Fourier encoding. By simulation, this approach is shown to allow for acceleration of the imaging process under ideal circumstances by nearly three orders of magnitude.

Key words: arterial spin labeling (ASL); functional magnetic resonance imaging (fMRI); functional connectivity magnetic resonance imaging (fcMRI); image reconstruction; pulse sequence design

Introduction

FOR OVER 30 YEARS, magnetic resonance imaging (MRI) has been built upon the principles of Fourier encoding, and the reconstruction of uniform grids of voxels from Fourier-encoded data. For medical diagnosis, images are typically evaluated by the human eye, and the generation of images with high and uniform spatial resolution is sensible. However, for functional imaging of the brain, image data are rarely evaluated by eye, and are instead analyzed using a wide variety of computer-based methods to extract functional signals and parameters of interest from these voxel data. Data-driven methods, such as independent component analysis (ICA) (Smith et al., 2012; Tian et al., 2013), and related methods detect features in the data that give clues as to the functional organization of the brain, while strongly model-driven approaches, such as dynamic causal modeling (Friston et al., 2003), attempt to extract parameters of well-defined models from the data. Other approaches, such as functional parcellation (Craddock et al., 2012; Shen et al., 2013) and graph theoretic approaches (Bullmore and Sporns, 2009; van den Heuvel and Sporns, 2011; Zuo et al., 2012), seek to characterize the relationships among brain parcels or nodes.

To date the imaging and processing sides of the functional MRI (fMRI) workflow have been largely independent, and the point of communication between these portions of the process is the four-dimensional data set (three spatial dimensions plus time) that is the typical output of the image acquisition and reconstruction components. The size of this data

set is enormous, typically in the hundreds of millions of data points, and yet, for most of the common processing procedures (outlined earlier), this data is typically reduced by approximately three *orders of magnitude* by the estimation of the blood-oxygenation-level-dependent (BOLD) signal in a few hundred parcels, nodes, or networks.

This implies that the relevant information content present in the fMRI signal is typically three orders of magnitude smaller than the voxel count. This 1000-fold excess in data presents an opportunity for dramatic increases in both imaging speed and SNR efficiency. The high spatial resolution typically acquired in fMRI is largely used to make sure that data are not assigned to the wrong side of a sulcus, rather than because the resolution matches the size of the features of interest.

We point out here that the intermediate step of reconstructing voxels is not explicitly necessary, and postulate that the rate of information extraction from the brain using fMRI and other functional imaging methods can be dramatically accelerated by estimating the parameters of the functional networks/nodes directly, bypassing the reconstruction of voxels.

Theory

The Human Connectome Project (HCP) (Smith, 2013; Van Essen et al., 2012, 2013) is a good benchmark for state-of-the-art fMRI methodology, as the imaging and processing protocols that are emerging from that project are becoming *de facto* standards in the field. The HCP has pushed MRI technology firmly up against some of the current

limits on imaging speed and resolution. It has led the field in advanced gradient technology to move through k -space as fast as possible (Ugurbil et al., 2013), and has also pioneered parallel acceleration methods, such as simultaneous multislice imaging with blipped controlled aliasing in parallel imaging (CAIPI) (Eichner et al., 2013; Setsompop et al., 2012), to maximize imaging speed. For BOLD fMRI, whole-brain imaging at 2-mm isotropic resolution with subsecond temporal resolution is now the standard (Ugurbil et al., 2013).

In a typical 2-mm isotropic whole-brain scan, $\sim 500,000$ imaging voxels are reconstructed at each time point. However, from each of these imaging volumes, the information extracted for most applications is contained in the coefficients of <500 brain nodes or networks. So, why is so much data collected?

One simple answer is that more data provide more degrees of freedom, and the ability to estimate more parameters. This is almost always true in the time domain. More time points, if they are at least somewhat uncorrelated, provide more statistical power, generating higher sensitivity and/or shorter scan times. However, in the spatial domain, resolution comes at a steep price in SNR. For a given voxel volume V and total scan time T , $SNR \propto V\sqrt{T}$ (Edelstein et al., 1986). Here T is the total time that the analog to digital converters of the scanner are actively acquiring data; it is assumed that magnetization is prepared in a similar way, and the SNR refers to the signal-to-thermal noise ratio. Now, suppose that the features of interest in a particular scan are N times larger than V . If the scan was adjusted so that the voxel volume matches the features of interest, then the SNR per voxel would increase by a factor of N for a given total scan time. This is partially mitigated by the fact that for the higher-resolution scan, N voxels can be summed to estimate the signal from the feature of interest, but this only recovers a factor of $\sqrt{N}$ in SNR, as the summation process increases the signal by a factor of N , but also increases the noise by a factor of $\sqrt{N}$. Thus, the final SNR of the estimate of the signal in the feature of interest is higher by a factor of $\sqrt{N}$ in the lower-resolution scan. As noted previously, the features whose signals are ultimately estimated (parcels, networks, or nodes) are typically 1000 times larger than our 2-mm isotropic voxels, and therefore $N=1000$ in the above calculations. Thus, the SNR is theoretically up to $\sqrt{1000} \approx 30$ times higher in a scan with feature-matched resolution, and an increase in scan time by a factor of 1000 in the 2-mm-resolution scan would be required to recover this SNR by averaging!

A second reason that high spatial resolution can be useful is that the brain is a spatially complex structure, and 2-mm resolution is approximately what is required to reliably localize the signal from each voxel to the correct location in gray matter, without contamination across sulci. While this is a valid anatomical consideration, the enormous cost in SNR should be understood.

If the geometry of the networks or nodes present in a given subject is known, then the data acquisition and reconstruction process can be reformulated as a function of the signal in each network or node. The distribution of signal across these networks/nodes is referred to here as network/node maps or nMaps, to distinguish them from spatially uniform voxel maps (i.e., conventional images). Example structures of potential nMaps are shown in Figure 1. For parcels, the value of one element in the nMap is the BOLD signal at

FIG. 1. The nMap is a construct that describes the coefficients of parcels, nodes, or networks over time. Network images adapted from Malaak et al. (2012).

one time point, integrated over the spatial extent of the parcel or node. The nMap for a set of networks is slightly different in that it consists of coefficients or “weights” for the BOLD signal associated with each of many potentially spatially overlapping networks. For example, for networks identified by ICA, the n th row of the nMap would be the time course associated with the n th network component.

For a given k -space trajectory, the linear relationship between the nMap coefficients and the MR signal can be expressed as a simple encoding matrix that is in general not a Fourier matrix, but instead one derived directly from the MR signal equations integrated over the arbitrarily shaped networks or nodes. For a k -space trajectory (k) and nMap coefficients (x), the MR signal S can be expressed as follows:

$$S = Ax$$

where the elements of the encoding matrix A are

$$A_{ij} = \int_{V_j} w_j(r) e^{ik_i \cdot r} dV$$

Here the integral is over the volume V_j of the j th nMap element, the weight w_j is the relative weight of the element over space, and i is indexes over the k -space points acquired. For elements V_j that fall over a uniform Cartesian grid, A reduces to the Fourier encoding matrix.

Our goal is to estimate the elements of the nMap directly from undersampled MR data, and hypothesize that this can dramatically increase the SNR of the estimates of the activation waveforms.

FIG. 2. Conventional functional magnetic resonance imaging versus direct calculation of the nMap.

Simulation

An example of direct estimation of nMap data from undersampled k-space data is shown in Figure 2 and compared with conventional estimation. In this example, a 288-parcel Yale atlas (Shen et al., 2013) was sampled at 2-mm isotropic resolution and used as a template. A range of BOLD signal values was assigned to each parcel, and Gaussian noise was added at a level such that in the image domain the average contrast-to-noise ratio (CNR) in each 2-mm voxel was 1. For BOLD imaging, a CNR that is on the order of unity is typical. For the conventional estimation, images were reconstructed by Fourier transform, and signals were averaged across the voxels of each parcel, yielding an estimate of the mean signal in each parcel. These values correspond to the values in one column of the nMap.

For direct nMap estimation, a small subset of 1500 of the noisy k-space data points were selected, representing an undersampling factor of 500, as shown. For this simulation, an efficient three-dimensional (3D) gradient trajectory was designed with low curvature throughout, so that the trajectory can be traversed quickly when subject to gradient slew rate constraints, and with a relatively incoherent aliasing pattern. The trajectory is shown in Figure 3, and is composed of approximately semicircular arcs extending from pole to pole, and intersects the points of a Fibonacci spiral at the equator.

FIG. 3. Three-dimensional (3D) K-space trajectory that provides a smooth 3D path, and incoherent aliasing patterns. The trajectory consists of semicircular arcs extending from pole to pole, with a Fibonacci spiral pattern at the intersection with the equatorial plane.

In this example, direct estimation of the nMap from 0.2% of the fully sampled data resulted in an RMS error that was 76% higher than that of the conventional estimation. In our initial testing, we found that the direct estimate always has higher error than the conventional estimate, which is sensible since it uses a subset of the same data. However, this increased error would be compensated by three averages of the 500 $\times$ undersampled data, which still leaves an under-sampling factor of over 150 $\times$, with the same net error.

We also performed similar simulations at CNR values of 0.2 and 5.0 for the 2-mm data (Fig. 4), and found similar results; that is, the increase in SNR with direct estimation of the nMap was $\sim 12\times$, and the optimal size of the sampled patch of k-space corresponded to a point spread function (PSF) that roughly coincides with the parcel size.

We noted previously that when the voxel volume is N times smaller than the features of interest, there is a $\sqrt{N}$ penalty in SNR for a given total scan time. In this example, the effective imaging resolution of the undersampled data, as measured by the width of the PSF (the FT) of the sampling pattern, was about 20 mm, which is approximately the average size of the parcels. Thus, the undersampled data were approximately matched in resolution to the features of interest. The parcels had an average volume of ~ 500 of the 2-mm voxels, and the direct nMap method should therefore have an SNR advantage of approximately $\sqrt{500} \approx 22$. We noted earlier that the direct nMap method has a 150 $\times$ advantage in scan time over the fully sampled conventional scan. This translates into an SNR advantage of approximately $\sqrt{150} \approx 12$, which is large, but it is clearly < 22 . Residual inefficiency may come from the fact that the parcel size does not match the imaging PSF exactly, and that the encoding matrix is not perfectly conditioned. The condition number for the encoding matrix in the aforementioned example was 2.5.

FIG. 4. Relative performance of conventional reconstruction and nMap estimation at three levels of contrast-to-noise ratio (CNR).

Raw 2mm images Number of data points in 2mm data: 699,840			
CNR of 2mm images	0.2	1.0	5.0
RMSE (2mm images)	12.8%	2.5%	0.54%
Number of data points for nMap	2103	1517	1378
RMSE (nMap)	19.5%	4.4%	0.88%
SNR efficiency improvement	12x	12x	14x

Discussion

At this point in time, the core limitation in fMRI is still primarily SNR, rather than spatial or temporal resolution. Even for the relatively low resolution of several hundred networks or nodes, fMRI data are typically averaged not only across experiments, but often across subjects to provide statistically reliable information. With potential gains of two to three orders of magnitude in scan time for a given SNR, the direct-mapping approach may be able to enable a *transformative step from multisubject averaging to near real-time in a single subject*. This would in turn not only increase efficiency, but also allow for the study of much more complex and dynamic thought and behavior. Even within a single subject, a gain in SNR efficiency would allow for shorter scan times and/or the exploration of more or more subtle tasks. The nMap approach can be applied to either connectivity- or task-based experiments, limited primarily by the circumstances under which the network mappings used are expected to be valid.

The currency of higher SNR can also be spent to simultaneously acquire images with multiple forms of contrast. One possibility is to acquire multiple echoes, which can be used to separate BOLD from non-BOLD sources of contrast (Kundu et al., 2012, 2013). Another is to acquire BOLD and CBF contrast, which can be used to estimate CMRO₂ and provide a more quantitative measure of brain metabolism (Blockley et al., 2013; Davis et al., 1998; Griffeth et al., 2013; Simon et al., 2013).

In the aforementioned SNR analysis, the calculated noise is thermal, rather than physiological, and is therefore not spatially encoded. In practice, both classes of noise are often important, and indeed the physiological sources of noise, such as subject motion, and cardiac- and respiratory-related modulations, tend to be dominant at lower spatial resolution (higher SNR). Physiological noise should therefore dominate the noise in nMap imaging. However, the higher temporal resolution afforded by the nMapping approach will also enable the use of a temporal sampling rate that is much higher than the cardiac frequency, and should therefore also allow for more effective removal of these noise sources in postprocessing.

In these examples, estimation of the nMap coefficients from the data was trivial, both conceptually and computationally, as it involved a direct least squares solution to the overdetermined set of equations, and the size of the problem ($\sim 2000 \times 300$) was small enough that it could be solved using a simple pseudo-inverse of the encoding matrix in less than one second. We note that even with multichannel

RF coil acquisition, the problem is likely to be small enough that direct least squares solutions are easy to compute.

One important feature of the proposed framework is the dependence upon the size of the features of interest. In our example (Fig. 2), the nMap had 288 parcels, each composed of ~ 500 two-millimeter voxels, leading to a potential for acceleration by a factor of 150. This potential acceleration factor decreases with decreasing feature size, and vanishes when the feature size equals the nominal imaging resolution. However, we anticipate that significant increases in SNR are still available at relatively high resolution, and the resolution limits of this approach are important to characterize. It is also important to note that while a few hundred nodes or networks are at the low end of the resolution spectrum, it is still a larger number than most analysis methods are currently considering. Detailed understanding of the workings of several hundred components in the dynamic human brain in single subjects would already represent a dramatic breakthrough in the study of brain function, and high-quality data at this resolution are therefore highly desirable.

While the potential gains in SNR efficiency described here are large, the implementation of this approach will certainly introduce nonidealities that will reduce the achievable gains in real-world applications. The most obvious sources of these nonidealities are (1) the requirement of good estimates for the geometry of the parcels or networks, and (2) the presence of non-BOLD, non-Gaussian sources of signal fluctuations in the data, such as motion and pulsations. The development of methods to characterize and address these effects will likely be an important task in the initial implementations of the nMapping concept proposed here. To address motion, we note that motion parameters can be estimated from the data itself and used to correct the data as a first preprocessing step. The detection of at least rigid motion parameters can be done with very high accuracy using low k-space data, as evidenced by the low-resolution nature of many motion navigator schemes (Maclaren et al., 2013). After motion correction, we expect that the high temporal resolution nature of the data should be helpful in detecting cardiac- and respiratory-related waveforms with high accuracy for use as regressors in subsequent analysis.

Potential methods to define the geometry of the nMap

Parcellation from standard atlases. This is perhaps the most straightforward source for the definition of the nMap,

though it is also likely the least accurate in terms of identifying brain regions that are truly homogeneous and well delineated in an individual subject. An atlas was used in the aforementioned example, but was assumed to be perfect. A test of this approach would be to use real fully sampled human data, and compare the measured functional activity to that calculated using the nMap approach with a fixed atlas and undersampled data. Conversely, this might also be a good way to compare the validity of different atlases.

Networks from ICA. Spatial ICA has historically been one of the most popular methods for detecting resting-state networks from fMRI data. A straightforward application in the context of nMap imaging is to use ICA on fully sampled data sets to determine the network structure for an individual, and subsequently assume that structure in the construction of the nMap bases and collect high-speed, high SNR data under those constraints.

Anatomical constraints. The use of purely anatomical constraints from high-resolution anatomical scans is another relatively straightforward application of the proposed methodology. One potential implementation of interest may be the use of anatomy to constrain the signal to gray matter. It has been shown that the width of the 2D PSF of the BOLD signal along the cortical surface at 3T is ~ 4 mm (Parkes et al., 2005). In the HCP processing stream, the cortical surface and deep gray matter are described using $\sim 90,000$ “grayordinates” at 2-mm resolution. This implies that at 4-mm sampling along the cortical surface, there are $\sim 22,500$ parcels, each with dimensions $4 \times 4 \times 3$ mm³ (3 mm is the approximate cortical thickness). If combined with parallel acquisition, then there may be enough data in a single 3D scan to estimate this nMap, and the possibility of single-shot nMapping of the whole brain at the resolution of the BOLD PSF would make an exciting goal.

Adaptive geometry. The use of a priori networks or nodes in the construction of the nMap basis may result in a poor fit to the data due to inaccuracy of the a priori information. Because the system is overdetermined, it may be possible to allow the data to adjust the geometry of the networks/nodes. A simple means of doing this would be to parameterize the geometry, and iteratively adjust these parameters to improve the quality of the fit to the data.

In the aforementioned examples, only spatial sparsity is exploited, and no use is made of sparsity in the temporal domain, such as that generated by temporal correlations. However, it is known that fMRI data are in fact sparse in the spatial-temporal domain, as demonstrated in the approximate low-rank nature of fMRI data in the K-T domain (K-T FASTER) (Chiew et al., 2013). In this approach, no basis functions are chosen to describe either spatial or temporal patterns, but low rank of the K-T matrix is assumed, and used to complete undersampled data. The nMap, as defined earlier, which has dimensions of network/node index and time, may also be of low rank, and the use of matrix completion methods may allow for further acceleration.

Finally, we note that similar concepts could apply to constrained reconstruction of nMaps from positron emission to-

mography (PET), magnetoencephalography (MEG), and electroencephalography (EEG) data, with the idea that predefined networks might in some cases bring the source reconstruction problem in these methods from underdetermined to overdetermined regimes.

Conclusions

We have introduced here the theoretical framework for direct imaging of brain networks. This approach is shown to have the potential to provide an order of magnitude increase in network imaging SNR, and thereby provide access to detailed dynamic imaging of complex brain activity in single subjects.

Author Disclosure Statement

No competing financial interests exist.

References

- Blockley NP, Griffeth VE, Simon AB, Buxton RB. 2013. A review of calibrated blood oxygenation level-dependent (BOLD) methods for the measurement of task-induced changes in brain oxygen metabolism. *NMR Biomed* 26:987–1003.
- Bullmore E, Sporns O. 2009. Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci* 10:186–198.
- Chiew M, Smith SM, Koopmans PJ, Blumensath T, Miller K. 2013. k-t FASTER: a new method for the acceleration of resting state fMRI data acquisition. In: The 21st annual meeting of the ISMRM (International Society of Magnetic Resonance in Medicine)—Discovery, Innovation, and Application—Advancing MR for Improved Health, Milan.
- Craddock RC, James GA, Holtzheimer PE, 3rd, Hu XP, Mayberg HS. 2012. A whole brain fMRI atlas generated via spatially constrained spectral clustering. *Hum Brain Mapp* 33:1914–1928.
- Davis TL, Kwong KK, Weisskoff RM, Rosen BR. 1998. Calibrated functional MRI: mapping the dynamics of oxidative metabolism. *Proc Natl Acad Sci U S A* 95:1834–1839.
- Edelstein WA, Glover GH, Hardy CJ, Redington RW. 1986. The intrinsic signal-to-noise ratio in NMR imaging. *Magn Reson Med* 3:604–618.
- Eichner C, Jafari-Khouzani K, Cauley S, Bhat H, Polaskova P, Andronesi OC, Rapalino O, Turner R, Wald LL, Stufflebeam S, Setsompop K. 2014. Slice accelerated gradient-echo spin-echo dynamic susceptibility contrast imaging with blipped CAIPI for increased slice coverage. *Magn Reson Med* 72:770–778.
- Friston KJ, Harrison L, Penny W. 2003. Dynamic causal modeling. *Neuroimage* 19:1273–1302.
- Griffeth VE, Blockley NP, Simon AB, Buxton RB. 2013. A new functional MRI approach for investigating modulations of brain oxygen metabolism. *PLoS One* 8:e68122.
- Kundu P, Brenowitz ND, Voon V, Worbe Y, Vertes PE, Inati SJ, Saad ZS, Bandettini PA, Bullmore ET. 2013. Integrated strategy for improving functional connectivity mapping using multi-echo fMRI. *Proc Natl Acad Sci U S A* 110:16187–16192.
- Kundu P, Inati SJ, Evans JW, Luh WM, Bandettini PA. 2012. Differentiating BOLD and non-BOLD signals in fMRI time series using multi-echo EPI. *Neuroimage* 60:1759–1770.
- MacLaren J, Herbst M, Speck O, Zaitsev M. 2013. Prospective motion correction in brain imaging: a review. *Magn Reson Med* 69:621–636.

- Malaak et al. 2012. Resting state models. *PLoS One* 7:e44428.
- Parkes LM, Schwarzbach JV, Bouts AA, Deckers RH, Pullens P, Kerskens CM, Norris DG. 2005. Quantifying the spatial resolution of the gradient echo and spin echo BOLD response at 3 Tesla. *Magn Reson Med* 54:1465–1472.
- Setsompop K, Gagoski BA, Polimeni JR, Witzel T, Wedeen VJ, Wald LL. 2012. Blipped-controlled aliasing in parallel imaging for simultaneous multislice echo planar imaging with reduced g-factor penalty. *Magn Reson Med* 67:1210–1224.
- Shen X, Tokoglu F, Papademetris X, Constable RT. 2013. Group-wise whole-brain parcellation from resting-state fMRI data for network node identification. *Neuroimage* 82:403–415.
- Simon AB, Griffeth VE, Wong EC, Buxton RB. 2013. A novel method of combining blood oxygenation and blood flow sensitive magnetic resonance imaging techniques to measure the cerebral blood flow and oxygen metabolism responses to an unknown neural stimulus. *PLoS One* 8:e54816.
- Smith S. 2013. Introduction to the NeuroImage special issue “Mapping the Connectome”. *Neuroimage* 80:1.
- Smith SM, Miller KL, Moeller S, Xu J, Auerbach EJ, Woolrich MW, Beckmann CF, Jenkinson M, Andersson J, Glasser MF, Van Essen DC, Feinberg DA, Yacoub ES, Ugurbil K. 2012. Temporally-independent functional modes of spontaneous brain activity. *Proc Natl Acad Sci U S A* 109:3131–3136.
- Tian L, Kong Y, Ren J, Varoquaux G, Zang Y, Smith SM. 2013. Spatial vs. temporal features in ICA of resting-state fMRI—a quantitative and qualitative investigation in the context of response inhibition. *PLoS One* 8:e66572.
- Ugurbil K, Xu J, Auerbach EJ, Moeller S, Vu AT, Duarte-Carvajalino JM, Lenglet C, Wu X, Schmitter S, Van de Moortele PF, Strupp J, Sapiro G, De Martino F, Wang D, Harel N, Garwood M, Chen L, Feinberg DA, Smith SM, Miller KL, Sotiropoulos SN, Jbabdi S, Andersson JL, Behrens TE, Glasser MF, Van Essen DC, Yacoub E, Consortium WU-MH. 2013. Pushing spatial and temporal resolution for functional and diffusion MRI in the Human Connectome Project. *Neuroimage* 80:80–104.
- van den Heuvel MP, Sporns O. 2011. Rich-club organization of the human connectome. *J Neurosci* 31:15775–15786.
- Van Essen DC, Smith SM, Barch DM, Behrens TE, Yacoub E, Ugurbil K, Consortium WU-MH. 2013. The WU-Minn Human Connectome Project: an overview. *Neuroimage* 80: 62–79.
- Van Essen DC, Ugurbil K, Auerbach E, Barch D, Behrens TE, Bucholz R, Chang A, Chen L, Corbetta M, Curtiss SW, Della Penna S, Feinberg D, Glasser MF, Harel N, Heath AC, Larson-Prior L, Marcus D, Michalareas G, Moeller S, Oostenveld R, Petersen SE, Prior F, Schlaggar BL, Smith SM, Snyder AZ, Xu J, Yacoub E, Consortium WU-MH. 2012. The Human Connectome Project: a data acquisition perspective. *Neuroimage* 62:2222–2231.
- Zuo XN, Ehmke R, Mennes M, Imperati D, Castellanos FX, Sporns O, Milham MP. 2012. Network centrality in the human functional connectome. *Cereb Cortex* 22:1862–1875.

Address correspondence to:

Eric C. Wong

UCSD Center for Functional MRI

9500 Gilman Drive Dept. 0677

La Jolla, CA 92093-0677

E-mail: ecwong@ucsd.edu