Toward a Biomimetic, Bidirectional, Brain Machine Interface

Andrew H. Fagg, *Member, IEEE* , Nicholas G. Hatsopoulos, Brian M. London, Jacob Reimer, Sara A. Solla, Di Wang, and Lee E. Miller, *Member, IEEE*

***Abstract*—The interest in Brain Machine Interface (BMI) systems has increased tremendously in recent times; many groups have become involved in this type of research, and progress has been quite encouraging. However, two fundamental limitations remain: 1) With a few notable exceptions, BMIs extract only kinematic information from the brain, ignoring the wealth of force or kinetic information also present in the primary motor cortex, and 2) most existing BMIs depend exclusively on natural vision to guide movement, lacking the rapid proprioceptive feedback that is critical for normal movement. The work reported here describes our efforts to address both of these limitations.**

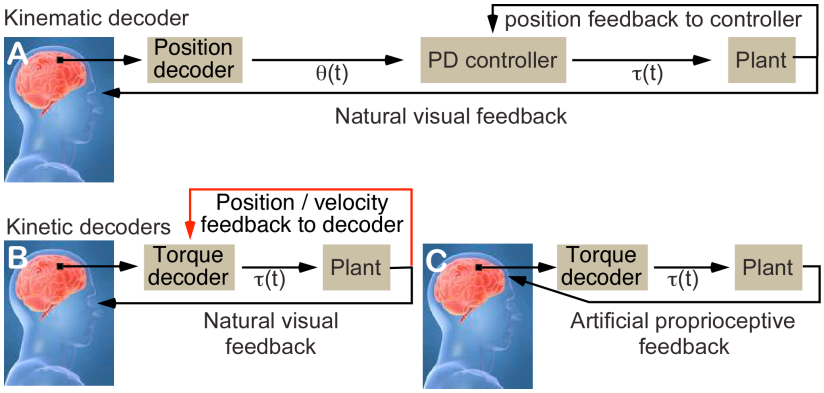

Fig. 1. A. Most brain-machine interfaces decode kinematic signals, such as joint angular position $\theta(t)$, from the activity of an ensemble of cortical neurons. The decoded signal is sent to a PD controller parameterized by a stiffness K and a viscosity B. The PD controller generates a torque $\tau(t)$ to drive the plant (limb) based on the error between the decoded position and the actual position feedback. B. In contrast, a kinetic decoder generates a torque $\tau(t)$ directly from the neuronal activity. The torque decoder can be augmented by delayed information about limb state (red pathway) C. In A and B, visual feedback regarding limb state allows the subject to control the device. Ultimately, we seek to provide proprioceptive feedback about limb state through microstimulation of the somatosensory cortex.

I. Decoding of Kinetic Signals for Limb Control

BMI-based control of a computer cursor or a robotic or simulated limb by monkeys has been demonstrated by several groups [1-5]. Up to now, the most common approach has been to extract only kinematic information from primary motor cortex (M1) When controlling an actual robotic limb, the production of the robot control signals (in the form of joint torques) is typically relegated to a feedback controller. This is typically accomplished with a PD (Proportional-Derivative) controller, which compares the "intended" state variables (position and/or velocity) to the current state of the limb (Fig. 1A). The output control signal for a single joint includes a component that is proportional to the error (the difference between the intended and actual state variables). Because an error must be present before the PD controller initiates a movement, the actual limb state will necessarily lag the intended state. The parameters of the PD controller can be tuned to minimize these delays under specific conditions, and to ensure that the arm will arrive at the decoded position without oscillation. However, this approach is limited because optimal parameter settings depend on the current dynamical context. For example, when the subject must grasp and apply a force to an object in order to push it, the change in system dynamics can result in suboptimal behavior of the PD controller.

One approach used in the robotics community to address this limitation of feedback PD control is that of feedforward, or inverted torque, control [6-9], in which an inverse dynamics model of the limb is used to estimate the appropriate torque signals given an intended trajectory. Instead, our approach is to assume that an inverse dynamics computation is performed within the central nervous system and that the results of this computation are reflected in neural activity and can be used as input to a torque decoder (Fig. 1B and C).

Our group has recently demonstrated that it is possible to decode both shoulder and elbow torque from M1 activity during a random target pursuit task [10]. In this task, the monkey produces arm movements that track a sequence of targets appearing in uniformly distributed locations on the screen, thereby sampling the position-velocity state space. We employ a linear filter decoding approach (Wiener filter) that considers the history of M1 neuronal activity as far as

Manuscript received April 23, 2009. This work was supported in part by NINDS grant #NS048845.

A.H.F. is an associate professor in the School of Computer Science, University of Oklahoma (fagg@cs.ou.edu).

N.G.H. is an associate professor in the Department of Organismal Biology and Anatomy, University of Chicago (nicho@uchicago.edu).

B.M.L. is a Ph.D. student in the Department of Physiology, Northwestern University (bml573@northwestern.edu).

J.R. is a Ph.D. student in the Department of Organismal Biology and Anatomy, University of Chicago (jreimer@uchicago.edu).

S.A.S. is a professor in the Departments of Physiology, and Physics and Astronomy, Northwestern University (solla@northwestern.edu).

D.W. is a Ph.D. student in the School of Computer Science, University of Oklahoma (di@cs.ou.edu).

L.E.M. is a professor in the Departments of Physiology, Physical Medicine and Rehabilitation, and Biomedical Engineering, Northwestern University (lm@northwestern.edu; Phone: 312-503-8677; FAX: 312-503-5101).

one second into the past in order to predict (decode) joint torque or Cartesian hand position or velocity.

Fig. 2 shows the reconstructed shoulder torque for a single trial that was not part of the data set used to train the decoder. The torque estimated from the limb inverse dynamics equations (black trace) was compared to the torque reconstructed from M1 discharge (blue trace). In general, the accuracy of offline torque reconstruction was nearly equal to that of Cartesian hand position and velocity reconstruction [10]. This result is perhaps not surprising given the relationship of these variables through the equations of motion and limb kinematics. However, it is important to note that the relation between torque and motion is nonlinear and that joint torque has a more complex temporal structure.

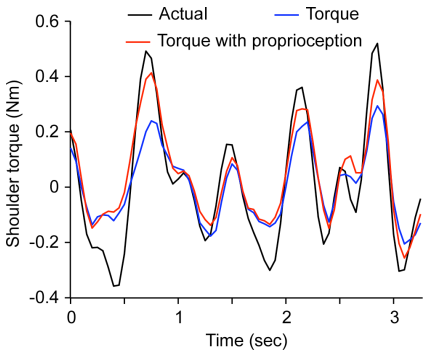

Fig. 2. Actual shoulder torque estimated from inverse dynamics (black trace) is compared to the torque decoded from M1 activity (blue trace). An improved torque reconstruction (red trace) was obtained by augmenting the decoder with delayed information about the actual joint positions and velocities.

We augmented the torque decoder with four additional inputs: the position and velocity of shoulder and elbow joints 100 ms prior to the prediction (Fig. 1B; red pathway). This decoder (red trace in Fig. 2) performed substantially better than the unaugmented decoder (blue trace in Fig. 2). Performance was a function of delay, and was maximal at the 100 ms delay [10]. Note that in this particular analysis, the limb state information corresponded to the monkey's actual limb movements; such information would not be available in a realistic real-time implementation.

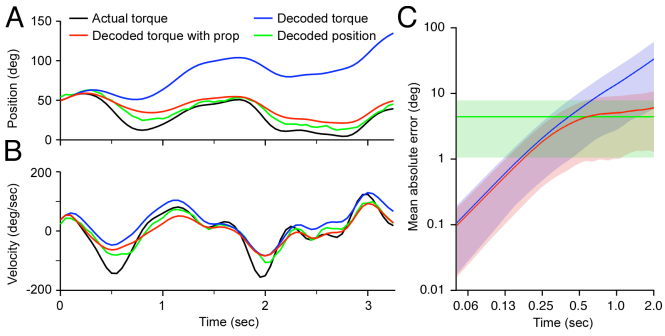

Fig. 3. Reconstruction of shoulder position (A) and velocity (B) based on kinematic and kinetic decoders. Reconstructions are shown as a function of time during a single trial. Each panel shows actual state variables (black), reconstructions based on position decoding of M1 activity (green), reconstructions based on torque decoding acting on a simulated limb model (blue), and reconstructions based on augmented torque decoding that incorporates delayed information about the simulated joint position and velocity, acting on the same limb model (red). C. Log-log plot of mean shoulder position error and its standard deviation as a function of time. In each case, the results of 720 trials are aligned at t=0 and averaged. Position decoder (green), torque decoder (blue), and augmented torque decoder (red).

The use of decoded torque in the direct, online control of an exoskeleton or prosthetic arm is fundamentally different from the use of decoded position. Errors in decoded position at any given time are independent of position errors made in the past. However, errors in decoded torque can accumulate in time and lead to substantial drift in the "intended" position of the limb. A key question is whether this drift occurs within a behavioral time scale that allows for correction by the subject.

Fig. 3 shows offline decoding of position and velocity signals during a single trial using several decoding approaches. Panels A and B show the actual shoulder position and velocity (black traces). The state of a simulated limb was initialized to the true state of the monkey's limb at time t=0. In an actual implementation, this would correspond to the subject's feedback guided correction of limb position. As anticipated, when the simulated limb was driven by the decoded torque, small errors in decoded torque accumulated over time and caused the simulated state to diverge from the true state after several hundred milliseconds (blue traces).

Panels A and B also show the evolution of the simulated limb when driven by the augmented torque decoder that incorporates information about limb state (red traces). Unlike the example in Fig. 2, the position and velocity signals used by this decoder represented the delayed state of the simulated limb, a signal that would be available in a realistic setting. The incorporation of these signals constrained the long-term drift of the torque decoder and achieved position tracking performance comparable to that obtained from the direct estimation of joint position (green trace).

Figure 2C shows a log-log plot of the mean tracking error in shoulder position averaged over a set of 720 trials as a function of the time elapsed from the start of the trial. As in panels A and B, the state of the simulated limb was set to ground truth at time t=0. Mean and standard deviation of the error of the position decoder are constant (green). The mean error of the torque decoder (blue) was smaller than the mean error of the position decoder until 400 ms after the beginning of the trial. This time scale suggests that the rate of drift is well within the range that can be corrected even with relatively slow visual feedback. Note that the position error of the torque decoder augmented with simulated position and velocity feedback (red) did not increase significantly beyond that of the position decoder.

We believe that the improvement in performance shown by the augmented torque decoder over the original torque decoder is due to a combination of factors. The incorporation of position and velocity information allows the augmented decoder to learn a strategy that resembles a PD controller, which can correct position and velocity errors. At the same time, because the decoder extracts torque information directly from the neural signals, it has the ability to combine kinematic information with information about the dynamics of the task.

These results show that under torque control, decoding errors accumulate slowly enough that the subject would have sufficient time to use visual feedback to correct the errors.

This suggests that online control by a kinetic torque decoder is feasible, and may offer advantages over simple kinematic decoding approaches under more complex and realistic conditions. We are currently evaluating the online performance of kinetic decoders.

II. Representation of Limb in Somatosensory Cortex

Existing efferent BMIs rely on visual feedback to guide movements and correct errors, functions normally dependent on the proprioceptive system and mediated by the fastest conducting nerves in the body. Patients suffering from loss of proprioception can move by relying on vision of their limbs, but their movements are typically slow, poorly coordinated, and require great concentration [11, 12]. We are pursuing the possibility of providing information about limb state to the subject via an afferent interface that operates in parallel with the efferent interface (Fig. 1C). The ultimate goal is to provide feedback about limb state by modulating the intensity of stimulation at one or more electrodes within primary somatosensory cortex (S1). Stimulus intensity would be determined by the state of the controlled limb transformed by an expression relating normal limb state to the discharge of S1 neurons. Previously we demonstrated that a monkey could perceive electrical stimulation of proprioceptive cortex (S1, area 3a) and discriminate between stimulation at different frequencies [13]. In addition to our own work, several other groups have demonstrated that a monkey can perceive stimulation provided through electrodes in S1 [14, 15], even, in one case, through the activation of single cells [16].

In order to characterize the patterns of neural activity in S1, we have used chronically implanted multielectrode arrays to record datasets containing 41 cells from one monkey (A) and 50 cells from another monkey (M) during the random target pursuit task described above. Of these 91 cells, 57 showed firing rate modulation related to the speed and direction of hand movement. We were able to reconstruct limb state from this ensemble activity using linear decoders similar to those used for the decoding of M1

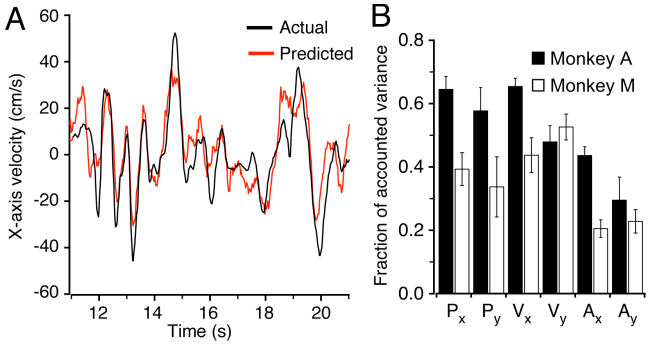

Fig. 4. A. Actual (black) and reconstructed (red) horizontal component of hand velocity as a function of time for an independent test sample. B. Fraction of accounted variance for X and Y components of position (P), velocity (V), and acceleration (A). Results are shown for two monkeys: A and M. Bars show mean +/- standard deviation for ten-fold cross-validated reconstructions: the data is divided into ten time segments or "folds", nine are used to fit the model, and the tenth is used to evaluate the quality of the model in cross validation. The process is then repeated ten times, using each fold as the test case.

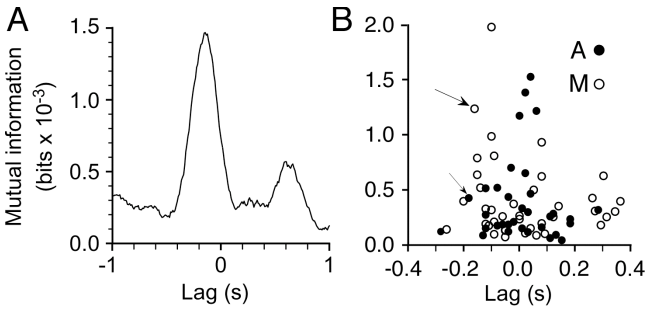

Fig. 5. A. Mutual information between cell activity and hand velocity for a single cell from monkey M (S1, area 2; top arrow in panel B). B. Empirical distribution of maximal mutual information versus the time lag of the peak for the 57 cells from both monkeys that showed firing rate modulation related to the speed and direction of hand movement. Arrows indicate cells shown in Fig. 7.

activity. Fig. 4A shows an example of velocity reconstruction along the x-axis. The linear filters used for extracting limb state information from S1 activity include lags of ±500 ms in order to include information from the neural discharge that lagged as well as led movement. Despite differences in the quality of reconstruction accuracy between the two monkeys (Fig. 4B), the greatest accuracy was for velocity and position reconstruction, comparable to that achieved when decoding position or velocity from M1 activity.

An alternative characterization of the information about limb state contained in S1 neural discharge follows from computing the mutual information between the discharge of each cell and the speed of hand movement. For the cell shown in Fig. 5A, the peak of mutual information is at -80 ms; this delay from limb state to S1 activity is expected from proprioceptive cells. The amount of velocity related information varied considerably across the recorded cells (Fig. 5B), but it should be noted that good kinematic reconstructions (as in Fig. 4) were not the result of only a few highly informative cells. Among all the S1, area 2 cells with significant (peak > 3 times noise) mutual information with the speed of hand motion, roughly half peaked at a negative time lag. The activity of cells whose mutual information peaked at positive time lags conveys information about future hand speeds. This effect could result from efference copy from motor structures, or from tuning to force or acceleration, both of which lead velocity.

In the random target pursuit task, it is not generally possible to identify the starting time of individual reaches, or to average over reaches in a specific direction. The advantage of this task is that it

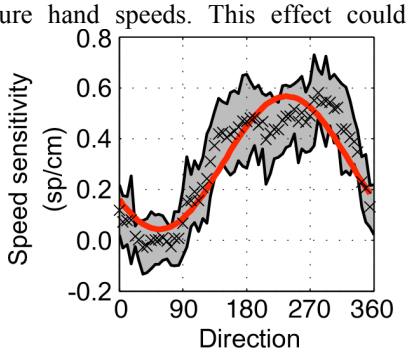

Fig. 6. Tuning curve of speed sensitivity for a single cell recorded from monkey A during the random target pursuit task. Most of the task-modulated cells from both monkeys had velocity-related discharge whose direction dependence was well fit by a sinusoid.

decorrelates position from its time derivatives; the disadvantage is that the computation of tuning curves is more complicated than for the more common center-out task. The time lag associated with the mutual information peak is used to pair up neural discharge with the projection of hand velocity onto any arbitrary axis, and thus relate the amount of discharge to speed for any direction of movement. A tuning curve can be computed by projecting along different axes (Fig. 6). Those recorded neurons that showed a significant peak in mutual information with hand velocity were also strongly modulated by both speed and direction of movement. The discharge of many cells was also modulated by position or acceleration, although the amplitude of modulation tended to be quite small.

The velocity modulation of the discharge reflected in the tuning curves suggests a simple model:

$$\lambda = m + k\|\vec{v}\|\cos(\theta - \theta_p) + s\|\vec{v}\| \ .$$

The firing rate λ is the sum of a baseline rate m, a cosine velocity term of amplitude k, where θ is the direction of movement and θ_p is the cell's preferred direction, and a direction independent speed term of amplitude s.

This simple cosine form provided a good description of the directional dependence of the discharge of many cells. The ability of this firing rate model to describe the observed discharge is illustrated in Fig. 7 for two cells, one from each monkey. Over the population of recorded cells, preferred directions exhibited a bimodal distribution, favoring those directions aligned with the monkey's forearm.

This relatively simple model of cell discharge could be used as the basis to stimulate S1 so as to mimic proprioceptive input from limb. It is important to note that this is only one possible model for conveying limb state information via S1 stimulation. More complex models, in which stimulus intensity is controlled by specific combinations of position and velocity, can be introduced if indicated by further analysis of the discharge data. Of course it must also be recognized that a stimulus train that perfectly matches the state-related activity modulation of a neuron will not lead to entirely normal neuronal discharge. The relation between stimulation and activation is unlikely to be linear. In any case, the stimulation will result in an abnormal, near synchronous activation of many cells in the vicinity of the electrode. It may be possible to monitor the resultant discharge of other neighboring cells as an indication of the effect of the stimulation. The nature of the evoked behavior of the monkey will serve as the most important indication of the efficacy of the stimulation. Finally, as in the efferent half of the bi-directional interface, we anticipate the need to rely on the brain's plasticity to adapt to the imperfect electrical stimulation.

III. Conclusion

Upper limb prostheses have gone through tremendous evolution in the past few decades, as it has been possible to replace mechanical control cables with modern electronics and motors. It is now routine to operate a 2 or 3 degrees of freedom prosthesis by sensing myoelectric signals from remaining proximal muscles through surface electrodes built into the socket of the device. However, despite the availability of increasingly complex multiarticulate arm and hand prostheses, many patients choose the traditional mechanically actuated prostheses, which provide a fairly natural feedback and simplify interactions with external objects.

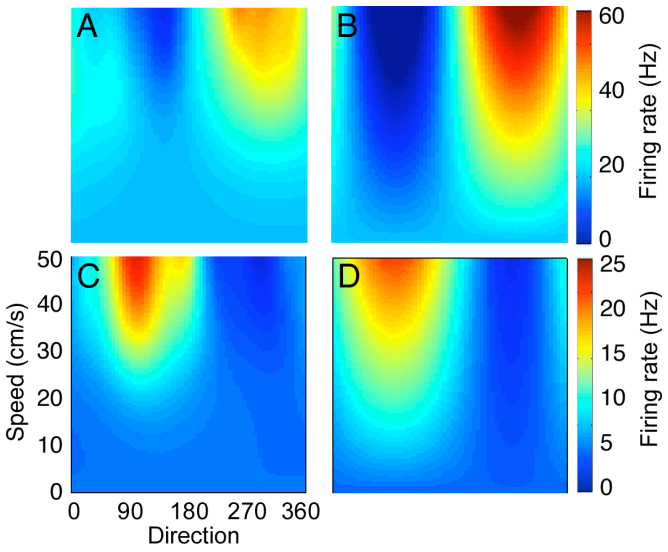

Fig. 7. A. Actual firing rate of a cell from monkey A, indicated by an arrow in Fig. 5B. The firing rate is shown by color code as a function of both speed and direction of movement. B. Firing rate of this cell as fit by the velocity sensitivity model. C. Actual firing rate of a cell from monkey M, indicated by an arrow in Fig. 5B. D. Firing rate of this cell as fit by the model.

The work reported here addresses two important aspects that need to be incorporated to the design of Brain Machine Interfaces in order to overcome these limitations. One is the possibility of controlling both force and position by utilizing information about both movement trajectory and the underlying muscle activity and forces provided by patterns of neural activity recorded in primary motor cortex. An improved efferent interface with the ability to control force in addition to position will allow the subjects to interact with an environment characterized by complex and changing dynamics. The second component is the incorporation of proprioceptive feedback through electrical stimulation in primary somatosensory cortex. An afferent mechanism for providing fast and effective proprioceptive feedback will reduce reliance on visual feedback and allow the subjects to experience percepts associated with natural movement.

References

[1] R. A. Andersen, S. Musallam, and B. Pesaran, "Selecting the signals for a brain-machine interface," *Curr Opin Neurobiol,* vol. 14, no. 6, pp. 720-6, Dec, 2004.

[2] N. Achtman, A. Afshar, G. Santhanam *et al.*, "Free-paced high-performance brain-computer

interfaces," *J Neural Eng,* vol. 4, no. 3, pp. 336-47, Sep, 2007.

[3] J. M. Carmena, M. A. Lebedev, R. E. Crist *et al.*, "Learning to Control a Brain-Machine Interface for Reaching and Grasping by Primates," *PLoS Biol,* vol. 1, no. 2, pp. 193-208, Nov, 2003.

[4] M. D. Serruya, N. G. Hatsopoulos, L. Paninski *et al.*, "Instant neural control of a movement signal," *Nature,* vol. 416, no. 6877, pp. 141-2, Mar 14, 2002.

[5] D. M. Taylor, S. I. Tillery, and A. B. Schwartz, "Direct cortical control of 3D neuroprosthetic devices," *Science,* vol. 296, no. 5574, pp. 1829-32, Jun 7, 2002.

[6] P. Khosla, and T. Kanade, "Experimental Evaluation of Nonlinear Feedback and Feedforward Control Schemes for Manipulators," *The International Journal of Robotics Research,* vol. 7, pp. 18-28, 1988.

[7] J. Craig, *Introduction to Robotics: Mechanics and Control*, 3 ed., Upper Saddle River: Pearson Prentice Hall, 2005.

[8] A. Liegeois, A. Fournier, and M. Aldon, "Model Reference Control of High Velocity Industrial Robots."

[9] M. Kawato, and D. Wolpert, "Internal models for motor control," *Novartis Found Symp,* vol. 218, pp. 291-304, 1998.

[10] A. H. Fagg, G. Ojakangas, L. E. Miller *et al.*, "Kinetic trajectory decoding using motor cortical ensembles," *IEEE Trans Neural Syst Rehabil Eng*, in press.

[11] C. Ghez, J. Gordon, and M. F. Ghilardi, "Impairments of reaching movements in patients without proprioception. II. Effects of visual information on accuracy," *J.Neurophysiol.,* vol. 73, pp. 361-372, 1995.

[12] R. L. Sainburg, M. F. Ghilardi, H. Poizner *et al.*, "Control of limb dynamics in normal subjects and patients without proprioception," *J.Neurophysiol.,* vol. 73, pp. 820-835, 1995.

[13] B. M. London, L. R. Jordan, C. R. Jackson *et al.*, "Electrical stimulation of the proprioceptive cortex (area 3a) used to instruct a behaving monkey," *IEEE Trans Neural Syst Rehabil Eng,* vol. 16, no. 1, pp. 32-6, Feb, 2008.

[14] R. Romo, A. Hernandez, A. Zainos *et al.*, "Somatosensory discrimination based on cortical microstimulation," *Nature,* vol. 392, no. 6674, pp. 387-90, Mar 26, 1998.

[15] N. A. Fitzsimmons, W. Drake, T. L. Hanson *et al.*, "Primate Reaching Cued by Multichannel Spatiotemporal Cortical Microstimulation," *J Neurosci,* vol. 27, no. 21, pp. 5593-5602, May 23, 2007.

[16] A. R. Houweling, and M. Brecht, "Behavioural report of single neuron stimulation in somatosensory cortex," *Nature,* vol. 451, no. 7174, pp. 65-8, Jan 3, 2008.