Neural Correlates of Uncertainty: A Dynamic Companion for ITS?

Ron Stevens [1] and Trysha Galloway[2]

[1] UCLA School of Medicine, Los Angeles, CA, 90230, USA
[2]The Learning Chameleon, Inc., Culver City, CA, 90230, USA

Abstract

This paper describes understandable neurodynamic models where the brainwaves of individuals and teams are analyzed in real-time, and measures are reported in terms of the frequency, magnitude and duration of uncertainty. These neurodynamic measures are causal intermediates between low level neural processes and the organizations that we recognize as being important for teamwork. They, are also ones that track closely with the hesitations and pauses we associate with uncertainty. Temporal and spatial brain region models reveal the brain regions, and possible causes of uncertainty. Machines can be trained to recognize the levels and the time-courses of the temporally and spatially defined uncertainty states, enabling them to support teams as active modelers, dynamic shapers and possibly oracles of future team behavior.

Keywords

Neurodynamics, Uncertainty, Machine Learning

1. Introduction

According to Sottilare and Hoehn [1] the process for identifying the states of teams and team members, and using them to select the appropriate adaptive instruction, generally consists of: Identifying behavioral markers to measure teamwork; Developing methods to automatically recognize unique behavioral markers; Associating behavioral markers with a unique teamwork state; and. Selecting and delivering an appropriate intervention based on the team member / team state.

Designing and implementing such systems becomes complicated with single-trial, ill- defined tasks. What actually is a learner model under these conditions and how can team and team member states be best defined and identified?

Conceptualizing and delivering this next generation of capabilities is challenging in that the dynamics of learning, retention and coordination during simulation training, or even making estimates of when that learning might be occurring, are poorly understood. Sometimes, the best that can be said is that learning likely occurred during the simulation, and perhaps more so during the debriefing. At the neuronal level however, it is increasingly apparent that learning is driven by unexpected events, i.e. those that cause uncertainty [2].

Uncertainty is a fundamental property of neural computation that we use to estimate the (perceived) state of our world. The brain draws from this uncertainty to access memories (the past) to imagine future possibilities and the actions needed to give the best outcomes, outcomes that might be orders of magnitude away in the future.

Humans maintain low levels of uncertainty by operating in familiar environments and situations where well-rehearsed sequences of cognition can be exploited. As a result, we think and act in terms of chunks of several seconds up to a minute, which help streamline the moment to moment activities [3.4]. To the extent that the planning and execution of these routines meet the immediate task requirements, the future will be predictable, and we avoid surprise.

Proceedings of the Challenges and Advances in Team Tutoring Workshop during the Artificial Intelligence in Education Conference, June 15th, 2021, Virtual from Utrecht, The Netherlands
EMAIL: immexr@gmail.com (A. 1); trysha@teamneurodynamics.com (A. 2).

Occasionally, unfamiliar environments or unexpected events increase our uncertainty about what to do next. When this happens the brain switches from exploiting past experiences to exploring new approaches [5, 6, 7]. In professional settings this exploratory uncertainty and the pauses and hesitations that it generates are often early indicators of deteriorating performance [8].

Currently it is difficult to predict how long uncertainty will last. The ability to rapidly and quantitatively measure uncertainty would have implications for educational and training efforts by supporting in-progress corrections, generating forecasts about future disruptions, or by using identified periods of uncertainty to target reflective discussions about past actions. The measures might also help target windows that are the most optimum for providing adaptive instruction.

The following sections outline a framework for developing increasingly understandable neurodynamic models that analyze the brainwaves of individuals and teams in real-time, and report the frequency, magnitude, and duration of neural correlates of uncertainty. These neurodynamic measures and models quantitatively scale from neurons to teams providing reports about performance levels and training gaps to stakeholders across the simulation community.

2. Modeling the Neurodynamics of Teams and Team Members

The challenge in developing performance-based evaluations using neural measures is not with the EEG measures themselves. Since the discovery of brainwaves, many measures have been developed using EEG, i.e. the frequency, amplitude (power), phase, complexity, scalp topology, ERPs, etc. An equally large number of methods have been developed for collecting, pre-processing and modeling the measures in real-time [9, 10, 11].

The challenge of broad-scale neural modeling has been that bottom up analytic approaches rapidly become complicated as most low-level neural processes are not in themselves directly causal to team performance but instead are the result of everyday cognitive activities that support seeing, listening, decision making, etc. It is when these activities are transiently amplified or modified by the context, that they assume greater importance for understanding teamwork.

Methods have been developed for estimating the neurodynamic correlates of uncertainty that are based on the information (not power or phase per se) in EEG rhythms [12-15]. The measure, Neurodynamic Information (*NI*), temporally bridges the gap between low level neural processes associated with everyday activities, and the organizations that we recognize as being important for teamwork. It is also a measure that tracks closely with the hesitations and pauses associated with uncertainty [8, 14, 16].

Detecting structure in data streams involves first deconstructing continuous data into discrete symbols and this requires deciding the number of partitions. Some EEG rhythms, like alpha waves (~10Hz), show either enhancing or suppressive neurodynamic properties depending on whether they are at a high or low power state [17] and so at its simplest, EEG amplitudes of a team member could be assigned any three symbols such that the states are easy to visualize and understand. In our studies activated states are assigned '3', deactivated states are assigned '-1' and neutral states are assigned '1'. The result is a data stream of 3's 1's and -1's.

The temporal structure (not power) can be estimated each second in this data stream by measuring the mix of the three symbols in a 60s segment that is updated each second as it slides over the data. If only one symbol was expressed in this 60s segment the entropy would be 0 bits; if there was an equal mix of the three symbols then the entropy would be 1.59 bits which is the maximum. So the fewer the symbols expressed in a window of 60s the more neurodynamically was the team member and the lower the entropy. Since it could be confusing having lower entropy mean higher organizations, the data plots are made more intuitive by calculating the Neurodynamic Information (*NI*) which is the bits of information when entropy values are subtracted from the maximum entropy for the number of unique symbols (Fig. 1A). So an entropy value of 1.0 would have a *NI* level of 1.585 minus 1.0 or .59 bits. As shown in Figure 1 B & C, dividing the signal into 6 or 12 states does not change the shape of the *NI* profile, just the level of information reported.

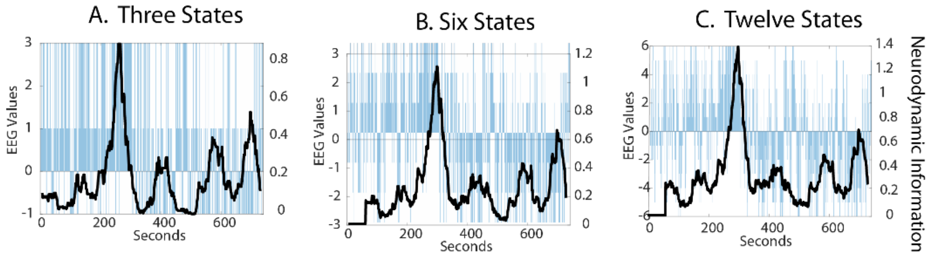

Figure 1: The EEG data stream of a member of an anesthesiology team member was divided into (A) three states, (B) six states, or (C) twelve states. The bar charts show the profiles of the EEG states while the solid lines trace the Neurodynamic Information.

A general, and perhaps more important point from this figure is that symbolically analyzing the structure of EEG amplitude creates a quantitative, and bounded scale of EEG organizations. It means that the neurodynamic information of any team of two persons who are performing any task where the EEG is separated into three levels will have *NI* levels between 0 and 3.17 bits, and for each team member the *NI* levels will lie between 0 and 1.585. These are values that can be quantitatively compared with other team performances, and can be aggregated for a class of trainees, or used to compare one training protocol to another. It means that the neurodynamic organization of one brain region can be directly compared with that of another brain region or across the frequencies of the 1-40 Hz EEG spectrum. Similar reasoning applies if the neurodynamic organization of a team is compared in the simulation scenario vs the debriefing, or across a critical healthcare event like intubation.

Neurodynamic information also contributes properties not always possessed by the amplitude or phase of brainwaves alone. For instance, neurodynamic information has been shown to link with the organization of team activities [12], or speech [18], or submarine navigation team expertise [19], or healthcare team expertise [14].

The emerging picture is that as simulations (and real-world events) evolve, the neurodynamic information accumulates and the bits accumulated are a function of the frequency, magnitude, and duration of periods of uncertainty. This feature applies to healthcare, military and pre-college teams and appears to be a general property of human performance.

The links between elevated *NI* and uncertainty were common during short periods (~ 1 min) of verbalized uncertainty [13], or during submarine navigation while the data needed to establish the submarine's position was being collected and shared among the navigation team [14]. The links have been extended to include medical students, hospital anesthesiologists and operating room staff (i.e. circulating nurse, scrub nurse and neurosurgeon etc.) during simulation training when they experienced difficulties ventilating a patient or deciding a course of patient management [16].

While originally elevated *NI* was described in the context of spoken uncertainty, elevations more generally occur during stressful periods whether or not someone was speaking. These associations were not unique to simulations as they were also seen in two neurosurgeons and an anesthesiologist during a live-patient surgery [20].

The sources of uncertainty observed parallel those described by Harencarova [21] for paramedics and include Recency Bias, Inadequate Understanding of the Situation, Technical Problems including equipment failure or a lack of understanding about how to operate the equipment, Inattention to Detail, and, Indecision regarding what to do next or choosing among a set of alternatives. The most frequent strategies used in response to these sources of uncertainty were trying to reduce the uncertainty and also by re-prioritizing actions.

3. Temporal and Spatial Brain Dynamics of Neurodynamic Information

Some of these properties are illustrated in this section where the neural signals of an anesthesiologist were quantitatively deconstructed across temporal and spatial brain scales and compared with task events during a training session. This session simulated a neurosurgery where after the patient was sedated the team was surprised by a rapid, and life-threatening, adverse systemic response to the anesthesia. The neurodynamics of the anesthesiologist are shown as she assumed team leadership under this situation. The dynamics began with a rapid

NI increase that corresponded with the team recognizing the difficulty and the formulation of an action plan (~800s). The large *NI* peak resulted from the scalp-wide contributions from half of the sensors, several of which (i.e. Fz, Pz, and T4) represented the involvement of the Default Mode Network (DMN). The DMN, has been proposed to be a network where prior information that is continuously accumulated over seconds to minutes, is melded with arriving extrinsic sensory information [22].

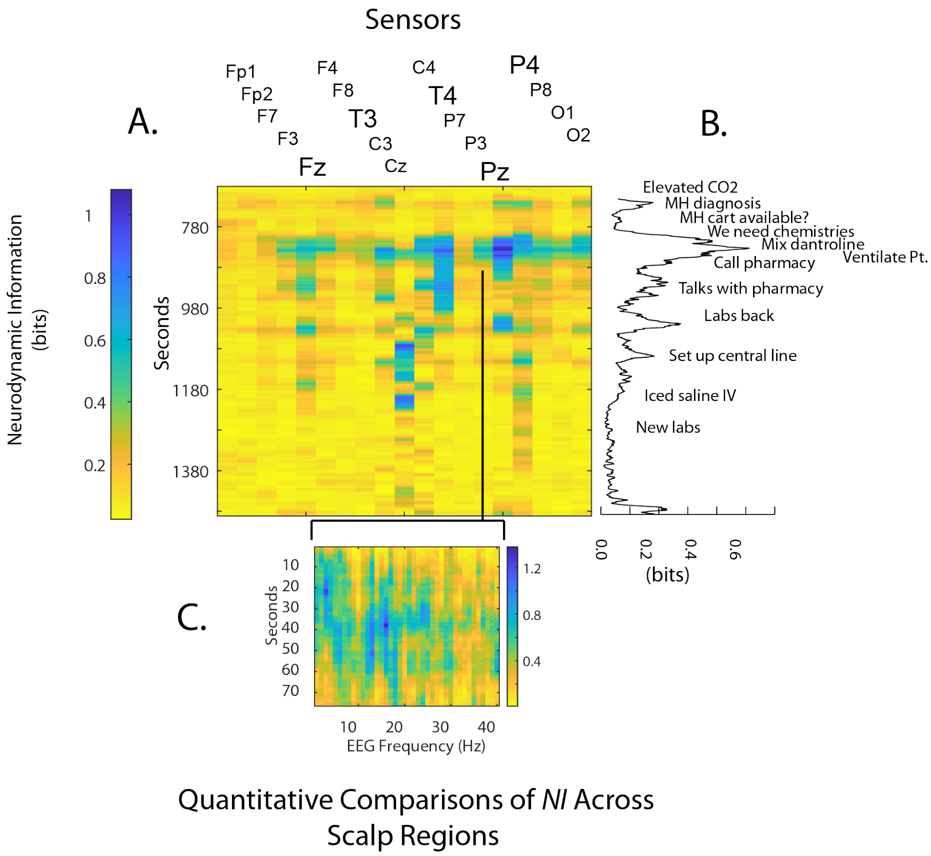

Figure 2: (A) The *NI* dynamics are plotted over time at each sensor of the Anesthesiologist. (B) The scalp-wide (i.e. averaged across sensors) *NI* dynamics are plotted along with events simultaneously occurring. **(C)** The elevated *NI* at the P3 sensor is expanded for one segment (800s-875s) and displayed across the 1-40 Hz EEG frequency spectrum; this example could be repeated for any sensor at any time if needed.

This temporally extended accumulated information is tied to the context of the stimulus, much like when reading a story the current chapter is continually linked with events and characters in the previous chapters. In this simulation, an absence of increased *NI* levels in DMN regions might, in fact, be a cause for concern.

Once the plan was formulated and the execution of the plan began, the *NI* decreased and the neurodynamics shifted from the multi-sensor DMN toward discrete peaks and more localized neurodynamic activity. As the patient stabilized the *NI* returned to near baseline levels.

The band of *NI* activity at the P3 sensor during the initial peak was further parsed across the 1-40 1Hz frequencies. The discrete nature of the peaks was preserved, first in the theta region (~4-5 Hz) and then in the beta region (~17 Hz), which may have implications for training machines as described below.

4. A Machine Learning Perspective of Neurodynamic Uncertainty

One of the hallmarks of human behavior is that of all the physiologic, mental, and spatial states we can be in, we often occupy and return to specific states; these support our routines [23]. Such repetition of behavioral states is particularly likely for teams performing tasks with repeating subtasks. During submarine navigation the position of the submarine is estimated every 3 min using multiple navigational aids [12]. The uncertainty associated with this process called 'Rounds' was investigated using self-organizing (SOM) artificial neural networks trained to recognize pattern variations in the *NI* peaks associated with verbalized uncertainty [13, 24]. These networks were trained to recognize sixteen uncertainty states with a topology ranging from no uncertainty to high levels of uncertainty. These networks represent uncertainty not only by the magnitude of uncertainty, but also by the profiles (Fig. 3).

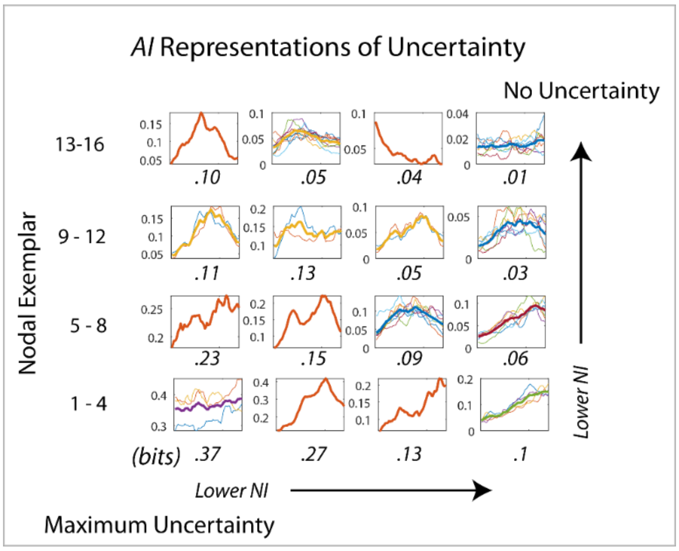

Figure 3: Topology of a Self-organizing Map for Uncertainty. The numbers below each exemplar indicate the average *NI* bits for the state across exemplars. The arrows indicate that uncertainty decreases from left to right as well as bottom to the top.

Transition maps can be used to identify states of uncertainty, their average and state-specific durations, as well as transition probabilities to future states (Fig. 4). The average durations between state changes for the four team members was 31.6 ± 40s (Figure 4A). This duration between state changes for the scalp-wide *NI* dynamics was similar to that calculated for frequency-specific duration using the Matlab® findpeaks.m function (26.3 ± 14s) or by averaging the durations between individual SOM state changes which was 25.4 ± 16s.

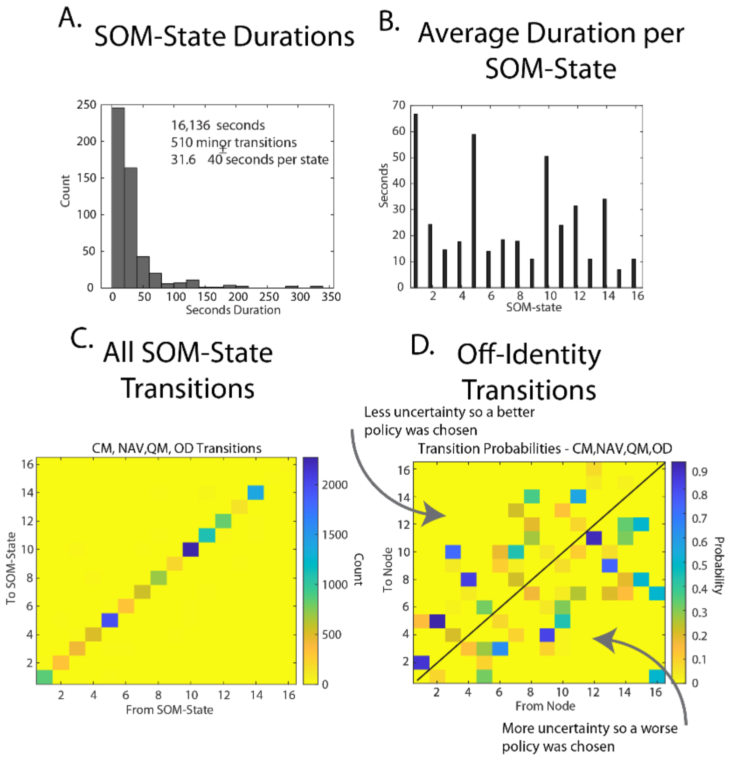

Figure 4: SOM-state durations and transitions. **(A)** Histogram of the durations of all SOM-states. **(B)** Average duration of each SOM-state is shown in black, the standard deviations in gray. **(C)** Transition matrix *From* a SOM-state (x axis) *To* the next state (y axis). **(D)** The minor state transitions were visualized by first removing the identity-line SOM-states. The color bar shows the probability of transiting to the next state. The identity line separates the transitions that will result in lower uncertainty (above the line) or higher uncertainty (below the line).

The durations at each SOM-state were variable both in terms of the average times (Fig. 4B, black bars) and standard deviations (Fig. 4B grey bars) with SOM states 1, 5, and 10. These states also represented upwards trajectory trends with few indications of an imminent decrease in uncertainty. The strong identity line in the transition matrix indicates that most SOM states were temporally persistent (Fig. 4C), i.e. once at a state there was a tendency to remain there.

There were also 510 SOM-state transitions that the team used to move along the identity line i.e. *From* state 6 *To* state 3 (Fig. 4D). The across-state transition probabilities were not evenly distributed across the state space but tended to divide themselves to those above and below the identity diagonal. With the topology shown in Fig. 3, those above the line represented transitions to a state of lower uncertainty. An example would be SOM state 7 where the most likely transitions led to a state of lower uncertainty, SOM states 8, 10 &13. Those below the identity line represent transitions to a higher state of uncertainty an example would be states 5 or 10.

These transition probabilities are meaningful in the sense that they could forecast a major future difficulty of the team on a second simulation [24].

5. Inserting Machines into the Tutoring Process

How can the dynamics of *NI*-related uncertainty be used for tutoring? First, the machine could function as a passive provider of post-hoc dynamic information. Training simulations are labor and time intensive, especially the after action reviews where much of the learning is thought to occur. The simple provision of *NI* maps like that in Figure 2 to the instructor before debriefing could help target areas for focused discussion. Alternatively they could be provided to trainees as supports for self-debriefings or discussions.

These ideas can be extended to training facilitators. The time required to achieve competency as a debriefing facilitator is variable, requiring the expertise of an experienced mentor. Factors such as understanding of key principles, practice opportunities, self-reflection and expert feedback are all important. Just as simulation and debriefing have the potential to standardize clinician education, a combined analysis of team *NI* and video, in conjunction with discussion and expert feedback, has the potential to standardize the formation of debriefers and decrease time to competency.

Moving past the level of debriefing, real-time monitoring of individual and team neurodynamics will also endow machines with an ability to understand the immediate changing state of human uncertainty allowing them to participate as active modelers, dynamic shapers and possibly oracles of future human behavior, in essence, providing continual answers to the question 'How is this team doing? [25].

In this regard, one of the more powerful advantages of using uncertainty-based artificial intelligence (*AI*) systems for education is that they will remember what a team, or team member struggled with in the past. Furthermore they will remember the magnitude and duration of these struggles as well as brain-locations and frequencies, and whether trainees would be likely to figure it out for themselves or require feedback. As performances accumulate and models are expanded and refined, the *AI* system could use its increased understanding of the behaviors of uncertainty to suggest new curricular designs to improve training efficiency and effectiveness.

Lastly, efforts to insert *AI* systems into organizations often fall short of expectations. Sometimes this is due to usability and trust issues, other times from current algorithms not being suited for ill-structured problems that are easy for people to perform, but hard for people to describe such as understanding or predicting the intent of other complex systems like teams. In education and training, these challenges can be compounded by the long-standing tradition of domain silos which can reduce the volume of the labeled exemplar data available for deep or reinforcement machine learning, further complicating the development of multi-purpose systems.

The very nature of uncertainty may facilitate acceptance of *AI* systems built around this construct, i.e. make it intelligible [26]. Humans spend much of their lives in states of uncertainty of various magnitudes and durations and so are familiar with the concept not only in principle, but also in practice. These understandings in conjunction with the bounded and quantitative *NI* scale (which is sensitive to training effects) will make it easier for users to accept recommendations from such transparent *AI* systems as they will better understand the factors considered in the recommendation.

6. References

[1] R. Sottilare, R. Hoehn, Sensing team interaction to enhance learning and performance during adaptive instruction. LNAI 2501, paper 3.

[2] C.J. Gillon, J.E. Pina, J.A. Lecoq, R. Ahmed et al., Learning from unexpected events in the neocortical microcircuit. bioRxiv 2021.01.15.426915;

[3] N.D. Daw, D.P. O'Doherty, Seymour, R.J. Dolan. Cortical substrates for exploratory decisions in humans. *Nature*, 441(7095) (2013): 876-879.

[4] D.W. Schneider, D. Logan, Chunking away task-switch costs: a test of the chunk-point hypothesis. *Psychol Bull Rev 22* (2015): 884–889.

[5] T.X. O'Rielly, Making predictions in a changing world: inference, uncertainty and learning. *Front. Neurosci 7*, 105 (2013).

[6] A. Soltani, A. Izquierdo, Adaptive learning under expected and unexpected uncertainty. *Nature Reviews Neuroscience. 20* (2019): 435-544.

[7] P. Domenech, S. Rheims, E. Koechlin, Neural mechanisms resolving exploitation-exploration dilemmas in the medial prefrontal cortex. *Science, 369* (2020): 1076-1079.

[8] M. Ott, A. Schwartz, M. Goldsmith, G. Bordage, L. Lingard, Resident hesitation in the operating room: Does uncertainty equal incompetence? *Surg Train Med Edu 52* (2018): 851-860.

[9] A. Delorme, T. Sejnowski, S. Makeig, Enhanced detection of artifacts in EEG data using higher-order statistics and independent component analysis. *Neuroimage, 34* (2007): 1443–1449.

[10] T.R. Mullen, C.A. Kothe, Y.M. Chi, A. Ojeda, T. Kerth, S. Makeig, T.P. Jung, G. Cauwenberghs, Real-Time Neuroimaging and Cognitive Monitoring Using Wearable Dry EEG. *IEEE transactions on bio-medical engineering*, *62*(11) (2015): 2553–2567.
[11] N. Bigdely-Shamlo, T. Mullen, C. Kothe, K.M. Su, K.A. Robbins, The PREP pipeline: standardized preprocessing for large-scale EEG analysis. *Front. Neuroinformatics* 9, B153 (2015).
[12] R. Stevens, T. Galloway, Are neurodynamic organizations a fundamental property of teamwork? *Front. Psychology, May, 2017*. doi: 10.3389/fpsyg.2017.00644 (2017).
[13] R. Stevens, T. Galloway, Teaching machines to recognize neurodynamic correlates of team and team member uncertainty. *Journal of Cognitive Engineering and Decision Making,* 13 (2019): 310-327.
[14] R. Stevens, T. Galloway, Parsing neurodynamic information streams to estimate the frequency, magnitude and duration of team uncertainty. *Front. Syst. NeuroSci.,* 01 February (2021).
[15] R. Stevens, T. Galloway, & A. Willemsen-Dunlap. Quantitative modeling of individual, shared and team neurodynamic information. *Human Factors, 60* (2018): 1022-1034.
[16] R. Stevens, T. Galloway, D. Halpin, A. Willemsen-Dunlap, Healthcare teams neurodynamically reorganize when resolving uncertainty. *Entropy, 18* (2016)*:* 427.
[17] W. Klimesch, Alpha-band oscillations, attention and controlled access to stored information. *Trends Cogn Sci. 16(12)* (2012): 606-617.
[18] J. Gorman, M. Martin, T. Dunbar, R. Stevens, T. Galloway, P. Amazeen, A. Linkens, Cross-level effects between neurophysiology and communication during team training. *Human Factors, 58,* (2015): 181-199.
[19] R. Stevens, T. Galloway, J. Lamb, R. Steed, C. Lamb, Linking team neurodynamic organizations with observational ratings of team performance. In *Innovative Assessment of Collaboration* (pp. 315-330). A. A. Von Davier, P. C. Kyllonen, & M. Zhu eds. Springer International Publishing, Cham, Switzerland (2017).
[20] R. Stevens, T.L. Galloway, A. Willemsen-Dunlap, Advancing our understandings of healthcare team dynamics from the simulation room to the operating room: A neurodynamic perspective. *Frontiers in Psychology, 10* (2019): 1660.
[21] H. Harencarova, Managing uncertainty in paramedics' decision making. *Journal of Cognitive Engineering and Decision Making, 11*, (1) (2018): 42-62.
[22] Y. Yishurun, M. Nguyen, U. Hasson, The default mode network: where the idiosyncratic self meets the shared social world. *Nat. Rev. Neuroscience, 22* (2021): 181-192.
[23] K. Friston, Am I self-conscious? (Or does self-organization entail self-consciousness?). *Front Psychol 2018,* Article 579. Available from://oi: 10.3389/fpsyg.2018.00679 (2018).
[24] R. Stevens, T. Galloway, Can machine learning be used to forecast the future uncertainty of military teams? Journal of Defense Modeling and Simulation, doi: 10.1177/1548512921999112 (2021).
[25] R.H. Stevens, J. Gorman, W. Zachary, J. Johnston, M. Dorneich, P. Foltz, How is this team doing, and why? In: R. Sottilare, A. Graesser, X. Hu, A. Sinatra. Design Recommendations for Intelligent Tutoring Systems: Volume 6 – Team Tutoring. Orlando, FL: US Army Combat Capabilities Development Command - Soldier Center (2018): 77-92.
[26] E. Vorm, A.D. Miller, Modeling user information needs to enable successful human-machine teams. In: D. D. Schmorrow and C. M. Fidopiastis (Eds.): HCII 2020, LNAI 12197 (2020): 445–65.