Cite as:
Roel Vertegaal, Timothy Merritt, Saul Greenberg, Aneesh P. Tarun, Zhen Li and Zafeiros Fountas. 2026.***Interactive Inference: A Neuromorphic Theory of Human-Computer Interaction.*** In J. A. Jacko (Ed.), The Human–Computer Interaction Handbook (4th ed.), CRC Press.

57 Interactive Inference
A Neuromorphic Theory of Human–Computer Interaction

Roel Vertegaal, Timothy Merritt, Saul Greenberg, Aneesh P. Tarun, Zhen Li, and Zafeirios Fountas

57.1 INTRODUCTION

There have been many efforts to develop psychological theories of human behaviour as it pertains to user interaction. Frameworks such as GOMS (Card et al., 1983) and ACT-R (Anderson et al., 1997) have been successful in allowing predictions on time performance when provided with a model user and a model interface. Similarly, models of human error (Reason, 1990) have been used to predict error as it relates to the design of interactive systems. However, there has not been a model that has comprehensively explained the relationship between human performance and error in a way that is generalisable. We believe developing such a model is important because it is a key to measuring cognitive load of users in real time. Current methods for measuring cognitive load fall short because they either rely on a questionnaire that can be difficult to administer during tasks (e.g., the NASA TLX, Hart & Staveland, 1988), or on physiological metrics that can be difficult to interpret (e.g., heart rate variability, electrodermal activity or pupillometry).

In this chapter, we discuss interactive inference, a theory of human performance as it relates to error that allows *direct measurement* of performance by first defining the signal-to-noise ratio (SNR) of a task. It incorporates the semantics of a task by modelling goals as predictions using Bayesian statistics that are solved using gradient descent. We set out to develop an approach to modelling tasks that examines both performance and error using a *relative* information-theoretical perspective. We will show how this approach might allow for wider comparisons of information processing capacity and error of users between otherwise incompatible tasks. While mathematically compatible with active inference, we note that our approach simplifies many of its key components and should be viewed as a practical approach rather than a strict implementation of the theory.

Key Takeaways

1. *Interactive Inference Model*: We introduce a simple neuromorphic framework called interactive inference that integrates human performance and error using information theoretical constructs. The model uses Bayesian statistics to interpret user goals as predictions, providing a way to measure performance by quantifying the Bayesian surprise processed per signal-to-noise ratio of tasks.
2. *Cognitive Load Measurement*: The model enables real-time estimation of cognitive load without relying on traditional self-reported questionnaires or physiological metrics. By modelling task complexity externally, it offers a more practical approach for evaluating user experience during interactions.
3. *Derivation of Human–Computer Interaction (HCI) laws*: Through known mathematical theorems and logic, we show how foundational principles of HCI – Hick's law, Fitts' law, and the power law of practice – emerge naturally from our framework through simple methods of gradient descent on a (Bayesian) surprise potential.
4. *Empirical Study*: To verify whether interactive inference could successfully model a task that does not yet have a known model, we demonstrate through an empirical study of a car-following task that user performance follows the logarithm of the SNR of the distance to a lead car.
5. *Research Agenda*: By discussing the limitations, benefits, and open research questions of this work, we invite further research to extend the theory and explore its practical limitations. This includes further empirical studies that test the validity of the theory as applied to different task-modelling scenarios.

57.2 BACKGROUND

57.2.1 Information Theory and Three Laws of HCI

Theories of the way the human brain processes information are as old as information theory itself. The very idea that humans might act like a computing circuit arrived shortly after the introduction of the microprocessor. Although most computing logic is deterministic, Shannon (1948) defined self-information as probabilistic: the lower the probability of a message, the higher its information content. One Shannon bit of information is defined as the negative binary logarithm of this probability (negative log probability):

$$I(x) = -\log_2\left(P(x)\right) \tag{57.1}$$

Shannon went on to define that the information capacity C of the transmission of a message through an analogue channel has a maximum that is the (binary) logarithm of

DOI: 10.1201/9781003401797-67

BK-TandF-JACKO_9781032513485-260130-Chp57.indd 1408 27/06/26 10:26 AM

the ratio between the signal strength S and the amount of Gaussian noise N in the channel, multiplied by some constant b, as expressed by the Shannon-Hartley Theorem:

$$C = b \cdot \log_2\left(\frac{S+N}{N}\right) = b \cdot \log_2\left(\frac{S}{N}+1\right) \quad (57.2)$$

where b is some empirically derived bandwidth parameter. This equation describes the capacity of information sent through the channel in terms of SNR. Note that without the negative sign, the SNR represents the reciprocal of the probability of a transmission. A high signal allows for more uncertainty to be carried, as does a lower noise. This means the very nature of information is statistical, with underlying equations based on (e.g., normal) distributions. This is important because humans, unlike traditional computer algorithms, operate in a probabilistic environment. That is, their exact behaviour cannot be computed; however, it can be modelled using (normal) probability distributions. Using this framework, Card et al. (1983) developed the model human processor (MHP), a comprehensive set of heuristics for predicting human performance in simple computer tasks. Statistical elements to these models were derived from empirical observations such as those performed by Hick (1952) and Fitts (1954). They demonstrated that human performance in certain tasks indeed seems to follow a similar pattern to the Shannon-Hartley Theorem: a binary logarithm of the SNR of the task. For example, Hick's law describes the response time RT for selecting from a number of alternatives n, plus 1:

$$RT = a + b \cdot log_2(n+1) \quad (57.3)$$

Here, a and b are some empirically derived constants. b models the time per unit of difficulty (the log component of the task). Response time is derived by multiplying the amount of information in the task by this empirical constant. Fitts (1954) derived a very similar law describing the movement time MT of a hand towards a target along a single dimensional trajectory A with a target width of W (specific equation from MacKenzie, 1992):

$$MT = a + b \cdot log_2\left(\frac{A}{W}+1\right) \quad (57.4)$$

One can think of the power of the signal in this equation as the amplitude (one-dimensional distance) to the target, and the noise as the width of the target. Again, a and b are constants that are empirically derived. b describes an average time per bit of difficulty of the task that when multiplied with the difficulty gives the movement time prediction. It is important to note the difference between the Shannon-Hartley Theorem and these equations: these equations predict time.

The third law was derived from Snoddy (1926) and describes the rate of learning of a model human. Time T to complete a task follows a power law as a function of the number of trials of that task x (Crossman, 1959):

$$T = a \cdot x^{-b} + c \quad (57.5)$$

Again, here, a, b and c are empirically determined constants.

57.2.2 Predictive Coding and the Bayesian Brain Hypothesis

One of the drawbacks of classic information theoretical approaches to human performance modelling is that classic information theory does not describe the presence of memory of messages transmitted through a channel, nor proactive anticipation of those messages. According to Shannon, a bit of information is simply the negative logarithm of the probability of a message (Equation 57.1). Only if a message is less probable, does it contain more information. The semantics of how the human interprets that message is not modelled, nor whether the human already knows the message. This is problematic because knowing the message, or anticipating the next message, reduces its information content to zero as this reduces the uncertainty about the message to zero.

Instead of framing human processing as a sequential computer processor, McClelland and Rumelhart (1981)'s parallel processing model described perception as an interaction of top-down (cognitive/semantic) and bottom-up (sensory) elements. Rao and Ballard (1999) further elucidated this approach by suggesting that the brain builds up statistical knowledge of its environment, i.e., incoming messages, over time, and that this model is used generatively, i.e., top-down, to *predict* the subsequent messages. The error between the prediction and the messages would lead to an update of the prediction.

Prediction is convenient for more than one reason: not only does it allow humans to plan actions in their environment based on relevant information, but it also serves to compress the information content of their environment by processing only those messages that are surprising, i.e., unpredicted or insufficiently predicted. Predictive coding theory explains how the human brain can work so efficiently and explains the mechanisms behind attention and scheduling of cognitive resources based on the amount of surprise in the predictions (Baldi & Itti, 2010).

According to Rao and Ballard (1999), Bayesian inference (Bayes, 1763) is a natural candidate for a generative statistical model. Within this context, Bayesian inference estimates the probability that a particular state s (representing a meaning or cause) exists, given an observation o. This is called the posterior probability $P(s|o)$, i.e., the chance of state s after, or given observation o. It does so by multiplying two probabilities: the likelihood $P(o|s)$ of the observation o given state s, and the *Prior* probability $P(s)$ of state s *before* the observation. The result is subsequently normalised by

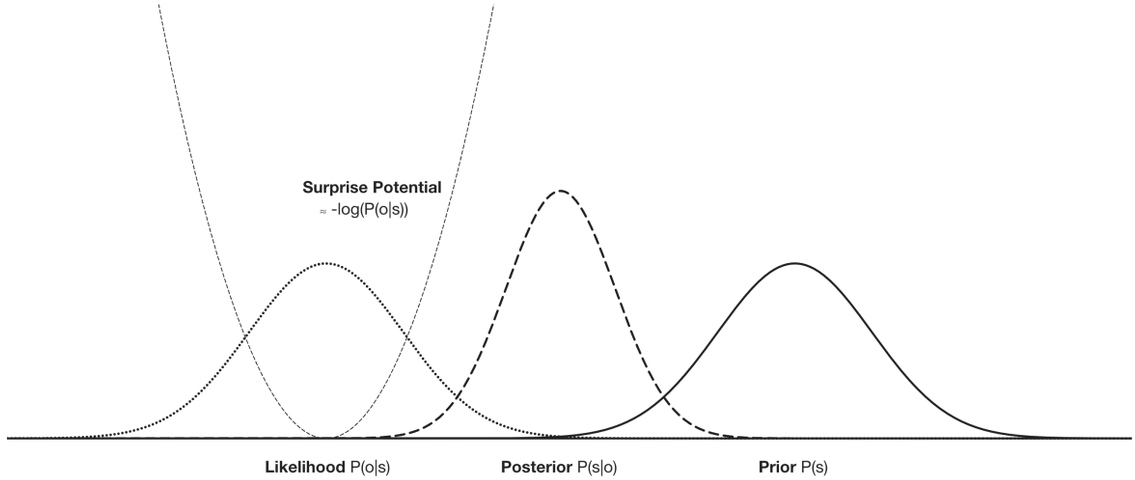


FIGURE 57.1 Bayes theorem. Distributions of the likelihood $P(o|s)$ of observations, left, given a prior $P(s)$ that describes the current state, right. In the middle is the calculated posterior $P(s|o)$ that describes the chance of the state given the observation. The squared function on the left shows an approximation of the negative log likelihood, the surprise potential that can be used to find the next posterior by gradient descent.

dividing by the overall probability of the observation $P(o)$ (see Figure 57.1):

$$P(s \mid o) = \frac{P(o \mid s) \cdot P(s)}{P(o)} \tag{57.6}$$

Unlike the frequentist approach to statistics, well known to the scientific community as hypothesis testing, Bayesian theory thus incorporates prior statistical knowledge in its testing of a hypothesis. The numerator serves as a generative model for predictions. The posterior $P(s|o)$ can be used as the subsequent prior for the next observation and as such, improve the prediction of future observations. This yields a different measure of uncertainty than Shannon information: the difference in information between the prior and posterior, or Bayesian surprise. Because Bayesian surprise incorporates the statistical relationship with prior knowledge it is semantic, describing the information content of the meaning of an observation under prior knowledge. Bayesian surprise is given by the Kullback-Leibler divergence (Kullback, 1997). It provides the relative entropy between the prior and posterior two distributions, or the information gained from one distribution to the other, in bits:

$$KL\big(P(s) \,||\, P(s \mid o)\big) = \int P(s) \cdot log_2 \left(\frac{P(s)}{P(s \mid o)} \right) ds \tag{57.7}$$

57.2.3 Entropy and Active Inference

Friston et al. (2006) further improved upon predictive coding theory by suggesting a mathematical framework for approximating the amount of Bayesian surprise (i.e., relative entropy) processed by the brain. In their framework, called active inference, learning and action are unified under a single principle: both use the difference between prior and (an estimate of the) posterior to make decisions. Minimising the difference informs when to learn or take action. When learning, the difference is minimised by updating the prior in light of the posterior (or, if necessary, by altering the likelihood), thereby updating the brain's generative model. When taking action, the difference is mainly minimised by altering the world, thereby updating the posterior to be closer to the prior knowledge of what the state of the world should be. Since we are concerned with the relative entropy between prior and posterior, the direction of updates does not principally matter.

Relation to Physical Entropy. To better understand the concept of relative entropy in information systems, let us use an analogy with the concept of entropy in the physical world. Physical entropy measures the logarithm of the number of microscopic states possible in a system (Boltzmann, 1866). Because particle interactions are often random, systems naturally evolve towards states with higher entropy, which can be loosely described as greater disorganisation. That is simply because high-entropy states are overwhelmingly more probable. This process is governed by the second law of thermodynamics (Clausius, 1879): in a closed system, entropy (analogous to disorganisation) increases over time and approaches equilibrium; it does not spontaneously decrease. Eventually, energy such as heat tends to become uniformly distributed throughout the system. Engines exploit free-energy gradients, which correspond to differences in entropy, to perform work. Random interactions also explain the natural tendency of heat to dissipate from your morning coffee into the environment. This illustrates the second law of thermodynamics: in closed systems, the entropy, which you can loosely think of as spread, increases.

Living organisms must counteract this natural tendency towards disorganisation by locally staying organised (Friston, 2012). That means they maintain their cells in some form of equilibrium called homeostasis. In return, they export entropy to their surroundings through waste products, including heat (Prigogine, 1977). They power this maintenance

using *physical* free energy imported from outside the system, from food that concentrates solar energy. Organisms must seek food while avoiding environments that may accelerate entropy increases. For this purpose, they evolved sensorimotor systems and neural circuitry that enable them to observe, learn, and act on their environment. In humans, these systems support homeostasis much like physical systems minimise physical free energy: they minimise the variational free energy (VFE) in their inferences, which bounds prediction error while avoiding needlessly complicated explanations (Friston, 2012). This means they learn to make correct predictions about their next state by reducing the Bayesian surprise in their predictions, for example, of where food might be. Bayesian surprise, then, is the relative entropy between the prior and posterior beliefs about the environment. It encodes the missing information between prediction prior to and after observation. The equations governing the statistical properties of particles and Bayesian surprise are thus closely related.

Action Mirrors Learning. The term variational refers to a complication in the calculation of Bayesian surprise between prior and posterior: that is estimating the overall probability of observations, the model evidence $P(o)$. For this reason and because the number of known states in the brain is extremely large, active inference uses a variational rather than an exact Bayesian approach to calculating the posterior (Beal, 2003). So, active inference uses a minimisation of VFE to reduce uncertainty about its predictions in a way that is mathematically analogous to the Feynman variational principle for physical free energy minimisation in statistical mechanics (Feynman, 1972). It minimises VFE because it is tractable, and from it an approximation of Bayesian surprise is obtained if the approximated posterior is close to the true posterior. In our interactive inference interpretation, we try instead to model Bayesian surprise more directly, as it is the actual measure of how much our beliefs change when we learn. When learning, sequential Bayesian updates of the prior allow humans to reduce subsequent Bayesian surprise. When acting, similar sequential Bayesian updates of movements occur (Körding & Wolpert, 2006). When learning, the prior is modelled as moving to the posterior; when acting, the posterior is modelled as moving to the prior. Once Bayesian surprise is minimal (or, rather, stationary), learning or action attenuates and the goal is reached. So, interactive inference differs from active inference in that we try to avoid approximating VFE, instead focusing directly on Bayesian surprise. To briefly summarise active inference before we continue, VFE is computed from what is tractable in Equation 57.6: the information in the likelihood function $P(o|s)$ and prior $P(s)$, given by the negative log joint probability of the generative model:

$$-log_2 P\left(o,s\right) = -log_2 P\left(o \mid s\right) - log_2 P\left(s\right) \qquad (57.8)$$

If distributions are normal, this produces a conveniently parabolic error function in which gradient descent (Courant, 1943) can be used to approximate the minimum free energy to subsequently approximate Bayesian surprise in a computationally efficient way, while avoiding explicit computation of the information contained in the intractable model evidence $P(o)$. This gradient descent is similar to that used for back propagation in neural networks (Rumelhart et al., 1986). Unlike classical neural networks, however, VFE allows unsupervised learning because it has an inherent reward or loss function: the minimisation of free energy, which does not require a separate ground truth. In terms of the user's experience: the reward is inherent when a system performs the function that the user predicted it to perform after some input. VFE is minimised, and Bayesian surprise is correspondingly reduced. In the remainder of this chapter, we will focus on how we might calculate Bayesian surprise more easily and more directly using the SNR of a task, while remaining within the general conceptual framework of active inference. We emphasise that interactive inference simplifies many aspects of active inference. Particularly, we do not implement model evidence terms or expected free energy calculations, focusing instead on a tractable framework suitable for HCI applications.

57.3 THE CONCEPT OF "TASK" AS A REDUCTION IN RELATIVE ENTROPY

Interactive inference implies that the core concept of HCI, the task, can be elegantly modelled as the minimising of Bayesian surprise between prior knowledge of a desired outcome (the prior or goal) and the posterior (the progress towards a goal given current observations).

57.3.1 Reducing Physical Relative Entropy

So we have a goal distribution, and a progress distribution, and the brain tries to reduce the relative entropy between them. Let us reconsider how, when we apply this concept to tasks in the real world, this can align with a statistical reduction in the *physical* entropy of the user's environment.

Let us consider a seemingly simple everyday task: organising socks in a drawer. Because the number of possible places where socks might end up other than the drawer is infinitely large, random interactions between, e.g., a human and the socks, make it much more likely that socks are anywhere but inside the drawer. The task, then, is to reduce the relative disorganisation in the location of the socks, or their physical entropy, relative to the location of the drawer. The goal distribution describes the probability of locations where socks need to be found, inside a drawer. The progress distribution initially describes the probability of the more random locations where socks may be found in the home. We are going to assume both are normal distributions. To execute this task and reduce the physical entropy in the organisation of the socks, we must expend physical energy (walk around and pick up socks, place in drawer), i.e., perform work.

According to interactive inference, the information used by the user's brain must follow a process proportional to the physical distributions: one that predicts any left-over physical relative entropy between socks and drawer. The brain is merely minimising the statistical difference in information about the location of socks and drawer that also defines their relative physical entropy, which is in and of itself a statistical property of objects. This means that the brain, in the process of acting to

minimise the relative physical entropy is also minimising the relative information entropy of the task, the Bayesian surprise. To do so, it calculates the statistical likelihood that an observation of a task object corresponds to the outcome of the task, after which it can estimate the chance that the task is complete using Bayesian inference (Friston et al., 2006).

57.3.2 Ten Postulates on Capacity and Error in Tasks

This leads to the first of ten postulates of interactive inference:

1. *Entropy Reduction*: The aim of any physical task is to minimise the relative physical entropy in the space of possible task outcomes. Conversely, any errors in the execution of the physical task lead to an increase in relative physical entropy. The processing that is required for the user's brain to execute a task leads to a reduction in relative *information* entropy, i.e., Bayesian surprise. In our example, this reduction is the *information* gain about the distribution of socks in the household. The aim of the task is to reduce the difference between where the user predicts the socks to be and where they actually are. When executing the task, the user's brain is computing the difference between the probability distribution of locations of socks outside the drawer (Likelihood), relative to the probability distribution of the goal: the available locations inside the drawer (Prior). In doing so it can, for example, estimate the time and effort required to complete the task, which leads to Postulates 2 and 3:
2. *Bayesian Surprise*: The relative *information* entropy that needs to be processed by a user to reduce the current error in the task to a satisfactory goal outcome is the Bayesian surprise. Bayesian surprise thus equals the amount of work that needs to be performed by the brain to complete the task given a current error, in bits. Bayesian surprise builds up over the scale of task outcomes to form a *surprise potential (see Figure 57.1).*
3. *Scale of Task Outcomes*: Relative entropy is measured along a scale of possible task outcomes *s*. This scale can be *any* feature of the task space that measures success or failure. In tasks in the physical world, relative location (i.e., distance) is often used as a scale of task outcomes. However, an example in the world of information might be a score on a test. This scale can be used to describe the (normalized) error in the task outcome.
4. *Surprise Potential and Difficulty*: The surprise potential of the task is an error function describing the relative information entropy or Bayesian surprise throughout the task, at any point along the scale of task outcomes. The task difficulty is the logarithmic coordinate of the surprise potential over task outcomes, that is, the binary logarithm of the normalized error. A fully accomplished task has followed the surprise potential to minimize the error. For example, the goal prior indicates that socks should be stored in a drawer. After the task of organising socks is complete, the probability of obtaining them in the drawer is 100%. This means the posterior distribution of socks now corresponds with the goal prior. By reducing both error and relative entropy along the surprise potential, the user has been able to make the current location of socks, the *posterior* or *progress distribution P*, correspond to the memorised *prior*, the *goal distribution G*.
5. *Goal and Progress Distributions*: Goal and progress of a task are defined as probability distributions on the scale of task outcomes *s*. The *goal* distribution *G* describes the prior: the probability of positive task outcomes on this scale. The *progress* distribution *P* describes the posterior: the probability of being in a current position along the scale of task outcomes *s*. Each distribution has a mean value μ on this scale. For the goal distribution, the mean μ_g signifies the value with the highest probability of a correct task outcome. For the progress distribution, the mean μ_p denotes the value with the highest probability of the current task outcome. Each distribution also has an associated variance. The standard deviation σ_g of the goal distribution defines the base noise scale of the goal. It can be scaled to fit the tolerances of particular instances of the task by a factor *z*. Thus, the actual acceptance region around μ_g is defined as $z\sigma_g$: completing outside this region constitutes an error. The variance σ_p^2 of the progress distribution describes the uncertainty with which the task outcome is perceived as moving toward the goal distribution; that is, the uncertainty in perceived progress on the task. The reciprocal of this variance, $1/\sigma_p^2$, is the precision with which current progress on the task outcome is controlled. In order for the goal and progress distributions to align, the user must not just reduce the distance between the means of the progress and goal distributions to within the *z*-scaled tolerance $z\sigma_g$, but also bring posterior uncertainty to within the *z*-scaled tolerance given for the task. A smaller *z*-scaled tolerance will require smaller updates on the scale of task outcomes as one progresses toward the goal. This is typically achieved not by changing the actual speed or precision $1/\sigma_g^2$ of the Bayesian update, but simply by following the differential of a surprise potential. The result is a halving of the difference between μ_p and μ_g in a given time period, defining one bit of update. A larger *z*-scaled tolerance allows faster completion time because users do not need to process the least significant bits resolving increasingly smaller steps near the end of the surprise potential. The ratio between the difference of means and the base noise scale reflects a normalized error. Although the task tolerance is defined by the *z*-scaled acceptance region $z\sigma_g$, successful completion requires the progress distribution

to become locally compatible with the goal distribution near the accepted region. That is, the posterior uncertainty need not equal the base variance of the goal distribution globally, but near the goal it must be of comparable scale. Under this local approximation, the effective variances of progress and goal may be treated as equal, allowing the negative log probability of the goal distribution *G* to be written as a quadratic surprise potential. This brings us to the next postulate.

6. *Signal and Noise*: We call the distance between the means of the goal and progress distributions on variable *s* the error *Signal* to the user's brain. It describes how far, and in which direction, the user needs to travel along the scale of task outcomes to reach the goal distribution to reduce the error to within the tolerance. The stronger the error signal, the more stimulus the task provides. However, this also means the more work has yet to be performed by the brain. The reciprocal of the error signal represents the *accuracy* of the task. The standard deviation of the goal distribution defines the task's base noise scale, specifying the base amount of variation permitted around the goal state. We call this the *Noise* parameter *(N)* of the task, as it relates to the maximum precision. Thus, the reciprocal of this Noise parameter represents a *precision*. *N* can be scaled with a parameter *z* to determine a variable threshold for tolerance within variations of the same task: *zN*. If we return to the example of organising your socks, the larger the average distance of socks, the *Signal (S)*, to the drawer, the harder it is to complete the task. The larger the *Noise (N)* scale within the drawer, and the further the drawer is opened *(zN)*, the easier it will be to complete the task. From this follows that the relative entropy of a task is a function of the ratio of the mean error *S* to the standard deviation *N* of the goal. This produces a normalized error. Note that this signal-to-noise ratio (SNR) is different from the one introduced by Shannon in that it describes a *relative error signal* to a goal:

$$\frac{\mu_p - \mu_g}{\sigma_g} = \frac{S}{N} \tag{57.9}$$

7. *KL Divergence*: Relative *information* entropy measures the reduction in Bayesian surprise, in bits, when the user tries to match progress and goal distributions. This relative entropy is calculated as a Kullback-Leibler divergence between the goal distribution *G* and progress distribution *P*, in bits. Mathematically, this is performed by first converting to bits, subtracting the two distributions (a fraction inside a log is a subtraction), then calculating the surface area using an integral. To be able to describe an empirically measured error, we include an empirical variable β in this equation that relates the KL divergence (where $\beta = 1$) to an observed empirical error rate (see Postulate 10):

$$KL\big(P(s)\,||\,G(s)\big) = \beta \cdot \int P(s) \cdot log_2\left(\frac{P(s)}{G(s)}\right) ds \tag{57.10}$$

8. *Equal Variance Condition*: The KL divergence can be greatly simplified in cases where the goal and progress distributions are normal, with identical variance. Bayesian surprise and thus difficulty can be calculated using a simple version of the KL divergence where we take the square of the SNR, where the signal is the distance between the means of the goal and progress distributions and the noise is the variance in the goal distribution. The mathematical proof of the relative entropy of two Gaussians with equal variance is, in fact, the motivation for the use of squared error in statistics and machine learning (Novak, 2015):

$$\frac{P(s)}{G(s)} = \frac{\sqrt{\sigma_g^2}}{\sqrt{\sigma_p^2}} \exp\left(-\frac{(s-\mu_p)^2}{2\sigma_p^2} + \frac{(s-\mu_g)^2}{2\sigma_g^2}\right) \text{and} \, \sigma_p = \sigma_g$$

$$\log_2 \frac{P(s)}{G(s)} = \log_2\left(\frac{\sigma_g}{\sigma_g}\right) + \log_2 \exp\left(-\frac{(s-\mu_p)^2}{2\sigma_g^2} + \frac{(s-\mu_g)^2}{2\sigma_g^2}\right) \tag{57.11}$$

The integral of Gaussian $P(s)$ multiplied with a function $log_2 \frac{P(s)}{G(s)}$ equals the expectation under $P(s)$:

$$KL\big(P(s)\,||\,G(s)\big) = \mathbb{E}_p\left[\frac{-(s-\mu_p)^2 + (s-\mu_g)^2}{2\ln(2)\sigma_g^2}\right] \tag{57.12}$$

The expectation under $P(s)$ equals:

$$\begin{aligned} \mathbb{E}_p\left[(s-a)^2\right] &= (\mu_p - a)^2 + \sigma_p^2 \\ \mathbb{E}_p\left[-(s-\mu_p)^2\right] &= -\left((\mu_p - \mu_p)^2 + \sigma_p^2\right) = -\sigma_p^2 \\ \mathbb{E}_p\left[(s-\mu_g)^2\right] &= (\mu_p - \mu_g)^2 + \sigma_p^2 \end{aligned} \tag{57.13}$$

Substituting expectations:

$$= \frac{(\mu_p - \mu_g)^2 + \sigma_p^2 - \sigma_p^2}{2\ln(2)\sigma_g^2} \tag{57.14}$$

We substitute the constant $\frac{1}{2ln(2)}$ with an empirical variable β that relates the KL divergence (where $\beta = \frac{1}{2ln(2)}$) to an observed empirical error rate:

$$KL\big(P(s)\,||\,G(s)\big) = \beta \cdot \left(\frac{\mu_p - \mu_g}{\sigma_g}\right)^2 = \beta \cdot \left(\frac{S}{N}\right)^2 \tag{57.15}$$

We note that the metric $\frac{\mu_p-\mu_g}{\sigma_g}$ is mathematically identical to Cohen's d (Cohen, 1988), a well-established statistical measure of effect size between two distributions. The resulting base equation is the canonical surprise potential. It is equivalent to the potential energy of a Hookean spring, where the error signal is the displacement from equilibrium and $1/N^2$ the stiffness constant (Feynman, 1972). Conveniently, it is also equivalent to the negative log of a Gaussian after dropping the normalization constant.

9. *Performance*: Human performance, or capacity to process relative entropy, follows the reduction in KL divergence as the user moves along the scale of task outcomes. This capacity to process Bayesian surprise, in bits, follows the Shannon-Hartley relationship (Equation 57.2), but with reduction in relative entropy, rather than absolute entropy. This equals the logarithm of our relative SNR plus 1.

$$C = b \cdot log_2\left(\frac{\mu_p - \mu_g}{z\sigma_g} + 1\right) = b \cdot log_2\left(\frac{S}{zN} + 1\right) \quad (57.16)$$

where b is some empirically derived constant describing the average performance. Human performance and error can thus both be described on the same scale of SNR. Note that this is a different scale than the task outcome scale. It describes the ratio of the error between means, and the standard deviation of the goal, i.e., the normalized error on the task outcome scale. Since performance is a logarithmic function of the SNR, while the amount of Bayesian surprise in a task is a squared function of SNR, users run out of capacity to process Bayesian surprise as the error SNR increases. The user's capacity to process Bayesian surprise can be measured, in real time, by calculating the average reduction in KL divergence per unit of base SNR. This capacity corresponds to the amount of cognitive load and/or sensorimotor load induced by a task because it describes the relative amount of Bayesian surprise processed by the brain: the information gained. As per Shannon-Hartley's theorem (Shannon, 1948), optimal performance is defined by the slope or average capacity of the logarithmic function on the SNR scale. Note that we can express the SNR scale in bits if we take its binary logarithm, in which case capacity becomes a linear function of binary log SNR.

10. *Error*: Human error equals the residual Bayesian surprise and is defined as the KL divergence ($\beta \cdot$ SNR squared) after remaining capacity ($b–b \cdot log_2(S/zN+1)$) becomes zero. Here, both β and b represent empirically determined constants. When capacity runs out, error cannot be kept arbitrarily low: it starts following an empirical KL divergence that increases with the square of the SNR. Error can thus be measured and predicted in bits in real time, using the same normalized error scale: that of SNR.

Note that error can also be modelled by calculating the negative log probability of the error rate per unit of SNR. Also note that the empirical error rate β in Equation 57.15 in that case is expressed in bits.

The ultimate purpose of any user interface design, then, is to allow the user to achieve Postulate 1 while avoiding Postulate 10. Together, our postulates give us a mathematical tool grounded in both physics and neuroscience that allows us to measure the capacity of the user as it relates to potential error, in real time. The easiest way to do this is to measure the mean distance to μ_g over some time window, dividing this by the standard deviation and squaring the result.

57.4 THREE LAWS OF HCI REVISITED

The three laws of HCI can be recast in terms of these earlier equations. These include Fitts' law, Hick's Law, the power law of practice and, potentially, others that are yet to be discovered using our method.

57.4.1 Fitts' Law

Körding and Wolpert (2006) and Wolpert (2007) strongly suggest that sensorimotor control is indeed guided by Bayesian prediction. Wu et al. (2006) and Williamson et al. (2022) also suggest that Fitts' law (Fitts, 1954; MacKenzie, 1992) is the result of the sensorimotor system performing Bayesian inference on movement. Figure 57.2 shows this process. A Fitts' law task involves moving the cursor (or hand) reciprocally between two targets or goals at distance A and click within width W. The goal can be modelled as a prior probability distribution $P(s)$. Width W sections off a tolerated area of this probability distribution that constitutes the visual target. This means W should not be interpreted as the variance of our goal distribution, but as an acceptance interval around the goal mean. If the base noise tolerance is $N = \sigma$,

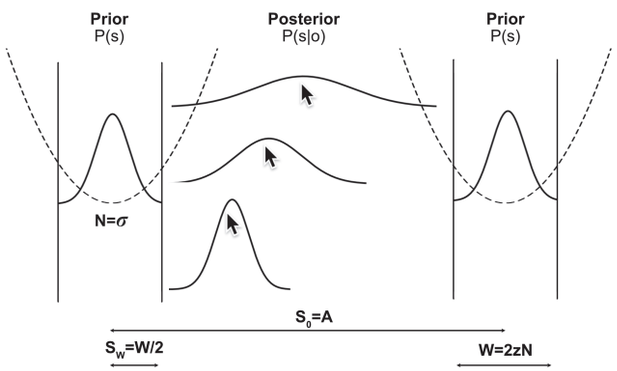


FIGURE 57.2 Bayesian Fitts' law. Distribution of the posterior probability $P(s|o)$ of the cursor being within the goal, from top to bottom: in the middle of a trial, at 2/3 s of the trial and at the end of a trial. Goal is the accepted W=2zN interval of the prior probability distribution $P(s)$, on both sides for a reciprocal trial. Cursor velocity follows a surprise potential (dashed), defined as the negative log probability of observing an error S under noise scale N.

the stopping threshold is a task-dependent multiple of this scale, $S_W = zN$, such that $W = 2zN$. In Fitts' Law experiments, both the distance A between the means of the priors and their tolerances, i.e., in the form of width W, are systematically varied (MacKenzie, 1992). The z-scaled prior distribution is evidenced over successive trials as an endpoint histogram (Welford, 1968). The figure also shows a cursor. The velocity of the cursor reflects uncertainty in the position of the cursor relative to the goal $P(s)$. This uncertainty is represented by the progress or posterior distribution $P(s \mid o)$, which must be estimated from observations of the cursor's distance to the goal given a tolerance. Figure 57.2 also shows how, as the cursor moves, velocity, and thus uncertainty in the posterior, increases to a maximum, after which it decreases to a minimal velocity or uncertainty once the target is reached. Near the goal, however, velocity, and thus uncertainty in the movement, must fall within the tolerance specified by the z-scaled prior so as not to click outside the target. In Interactive Inference, then, movement is modeled by a step-wise iterative update of the posterior toward the prior(s).

The size of this update is governed by the slope of a surprise potential placed on the center of the target(s). This surprise potential is approximately the negative log probability of the goal distribution, and equals the KL divergence in Equation 57.15: $1/2\ SNR^2$. The initial error signal S_0 is equal to A, while the noise N relates to a fixed precision estimate: N does not change with target size W. Rather, each W is evaluated as a threshold S_W on the same surprise potential, some multiple z of N, upon which movement control is attenuated. The calculation of each new estimate of cursor velocity is a Bayesian update, for which the brain uses the derivative, or slope, of the segment of the surprise potential delineated by S_0 and S_W ($W/2$). Indeed, Welford suggested early on that performance is a function of a signal-to-noise ratio in the brain and that there is a maximum capacity at which the brain can process a Fitts' Law task due to "*neural noise blurring or distorting signals*" (Welford, 1968). This describes the base noise scale N. MacKenzie (1992) suggested the following equation for describing the resulting movement time in a Fitts' Law task:

$$MT = a + b \cdot log_2\left(\frac{A}{W} + 1\right) \tag{57.17}$$

Here, MT equals movement time, and a and b are empirically determined constants. The equation is modelled after the Shannon-Hartley theorem in Equation 57.2. We can instead define a standard surprise potential U in the form of Equation 57.15 that can be minimized over time:

$$s = \frac{S}{N}, U = \frac{1}{2}s^2 \tag{57.18}$$

where the initial condition $S_0 = A$, and $1/N^2$ represents the maximum spatial motor precision $1/\sigma_g^2$ required for the cursor or hand movement in the task, thus defining an invariant geometric scale. This surprise potential has the same form as the potential energy of a spring. If we therefore express the movement as a first-order relaxation of this "spring" (i.e., derivative of Equation 57.18), we obtain:

$$\frac{\partial U}{\partial s} = s,\ \dot{s} = -\frac{1}{b}s \tag{57.19}$$

Here, $1/b$ governs the *temporal rate* or temporal precision at which the system relaxes along the surprise potential. It relates directly to the index of performance derived from Equation 57.17. The stop condition is when the cursor finds the near edge $S_W \leq zN$, where W equals $2zN$. The resulting behaviour corresponds to a binary search for the target in which the error in cursor position is controlled by halving the distance to the edge of the target, at each time interval, until the target edge is found. Each halving can be considered a *bit*. Equation 57.18 describes the control function that encodes the potential energy underlying the planning of the cursor (or hand) movement. Equation 57.19 models the corresponding force field and its first order relaxation over time. Modeling actual movement, however, requires converting this potential energy into kinetic energy. This can be achieved by adding a second order gradient descent that accounts for the input device's inertia (mass) and the muscles' viscoelastic stiffness (damping), which is beyond the scope of this chapter. In a reciprocal tapping task between the centres of targets, the relaxation of potential energy over time produces the logarithmic time result found in Equation 57.17 (Knuth, 1998):

$$S(t) = S_0 e^{-t/b} \Rightarrow MT = b \ln\left(\frac{S_0}{S_W}\right) = \frac{b}{\ln 2}\log_2\left(\frac{S_0}{S_W}\right) \tag{57.20}$$

Here, $S(t)$ is the time based error signal in this task, with S_0 the initial distance between the centre of the goals. S_W describes the stop condition in this task at half the goal distribution $W/2$. Empirical validation of this Bayesian underpinning of Fitts' law remains ongoing. However, Schmidt et al. (1979) showed the variance in motor impulse in a similar task to be linearly related to the variance in the resulting distribution. Our surprise formulation can be used to predict error, if error rates are converted to bits using negative log probability. It may also allow the comparison of capacity between tasks. This is important in cases where the overall capacity of the user's actions needs to be calculated in bits of unresolved Bayesian surprise. From our earlier discussion follows that once a user runs out of capacity, errors should track a KL divergence, some empirically determined surprise potential, after the user runs out of capacity:

$$E = \beta \cdot \left(\frac{S}{N}\right)^2 \tag{57.21}$$

Here, β is an empirically determined constant describing the error rate in bits (not time). There is some evidence for such error curve (de la Fuente, 2014). Welford (1968) demonstrated a method for post-hoc correction of goal distributions, so as to adjust for differences in error rates between different devices. His work appears consistent

with the above discussion. Wobbrock et al. (2008) also modelled the probability of an error, with good fit. Our model predicts that user error only tracks the KL divergence once capacity is reached, while their model describes the probability of error rather than the negative log probability. Guiard et al. (2011) analysed the error in different conditions of a Fitts' law experiment. Their RVE metric also appears to be consistent with Equation 57.21.

57.4.2 Hick-Hyman Law

Hick's Law, better known as Hick-Hyman law (Hick, 1952; Hyman, 1953) can also be derived from our surprise potential. Hick-Hyman law describes user response time when the task is to select one from a number of choices. Hick (1952) found that the response time was a logarithm of the number of choices $n + 1$:

$$RT = a + b \cdot log_2\left(n+1\right) \tag{57.22}$$

Here, a and b are some empirically derived time constants. b can be used to convert the difficulty of the task to a time estimate. Unlike Fitts' law, we model Hick-Hyman law as a learning task in which the user performs Bayesian inference to update a prior of possible targets with an observation of a single highlighted candidate to create a posterior of only one candidate. This means the roles of the prior and posterior distributions are not reversed: in this task we update the prior to the posterior. Hyman (1953) extended Hick's experiment to include unequal probabilities for each choice. He found that time results were a function of the entropy of the task:

$$\sum_{i=1}^{n} p_i \log_2\left(\frac{1}{p_i}\right) \tag{57.23}$$

where p_i is the probability of each choice in the set. Note that as the set gets smaller, the chance of a correct answer increases. As the set gets larger, it decreases by the number of elements in the set. This means the only relevant parameter in this law is p_i. According to Wu et al. (2018), evidence from MRI scans of the brain indeed suggests that the narrowing of the selection from the prior to the observation follows step-wise Bayesian posterior updates. According to them, the cognitive control network in the brain eliminates half the options at each Bayesian update, yielding a binary search with a logarithmic response time. Within our framework, Hick-Hyman law can be rewritten using the simple first-order gradient descent in Equation 57.19 on the canonical surprise potential from Equation 57.18, where, in the probabilistic view, S represents the total choice probability, equal to 1, and noise tolerance N denotes the probability of a single choice, equal to the reciprocal of the total number of admissible choices. The surprise potential (see Figure 57.1) represents an energy potential with higher energies (negative log probabilities) reflecting smaller probabilities. Competing neurons relax this potential over time, until one reaches the decision threshold S_N. Solving for time yields:

$$S(t) = S_0 e^{-t/b} \tag{57.24}$$

$$RT = b\ln\left(\frac{S_0}{S_N}\right) = \frac{b}{\ln 2}\log_2\left(\frac{S_0}{S_N}\right) = \frac{b}{\ln 2}\log_2(n+1) \tag{57.25}$$

where t denotes time, and, in the count ratio view, S_0 denotes the initial number of possible alternatives, including the case of no stimulus, $S_N = 1$ the stopping condition corresponding to the single resolved choice, and b the empirically determined time per decision bit.

57.4.3 Power Law of Practice

The third law of HCI, the power law of practice (Crossman, 1959), governs the learning curve that arises as users learn a new task. Again, this is a learning task in which the prior is updated to the posterior. The theoretical basis for the power law of practice was not explained: it is based on an empirical finding that was confirmed in many experiments (Crossman, 1959). These suggest that the amount of time T spent performing a trial of a task is a negative power function of the number of trials x performed:

$$T = a \cdot x^{-b} + c \tag{57.26}$$

where a and b are empirically derived constants. Within the context of interactive inference, this power law is explained as a reduction in the KL divergence between prior and posterior that occurs when users are learning a task. This reduction is a function of the SNR improving with every trial. In deep learning, this is known as minimising the negative log likelihood function as a function of SNR (Bishop & Bishop, 2023):

$$H = a - b \cdot log_2\left(\frac{S}{N}\right) \tag{57.27}$$

where a and b are empirically derived constants. This equation describes the reduction in entropy, in bits, as an improvement in the SNR of the task. The time taken per bit then provides the power law of practice.

57.5 THE CASE OF UNEQUAL VARIANCES

There are some cases in which the simplification of the KL divergence to a square function of the SNR does not work. Consider the example of the distribution of socks that need to be placed in the drawer. We note that we could consider this task a compound task in which we sum the individual KL divergences of the movement of each individual sock.

In this example, we are simply going to take their average, in an attempt to describe the overall relative entropy of the task with one KL metric.

We recognise that the surprise in this task is not just governed by the noise scale, size of the drawer or average distance of socks, but also by the number and distribution of socks in the room or house. This means that there maybe a discrepancy between the size of the distribution of socks, our posterior, and the goal, our prior. This discrepancy also increases the KL divergence describing the surprise of the task. To calculate the KL divergence, we need to capture this difference. While the original integral does this for any distribution, including non-Gaussians, it is computationally expensive and not formulated in terms of SNR. There exists a more general function that is mathematically identical to Equation 57.7, which describes the KL divergence in cases of Gaussian distributions with unequal variances (Moulin & Johnstone, 2014):

$$\log_2 \frac{P(s)}{G(s)} = \log_2\left(\frac{\sigma_g}{\sigma_p}\right) + \log_2 \exp\left(-\frac{(s-\mu_p)^2}{2\sigma_p^2} + \frac{(s-\mu_g)^2}{2\sigma_g^2}\right)$$

$$KL\big(P(s)\,\|\,G(s)\big) = \log_2\left(\frac{\sigma_g}{\sigma_p}\right) + \mathbb{E}_p\left[\frac{(s-\mu_g)^2}{2\ln(2)\sigma_g^2}\right] - \mathbb{E}_p\left[\frac{(s-\mu_p)^2}{2\ln(2)\sigma_p^2}\right] \tag{57.28}$$

The expectation under $P(s)$ equals:

$$\begin{aligned}
\mathbb{E}_p\left[(s-a)^2\right] &= (\mu_p - a)^2 + \sigma_p^2 \\
\mathbb{E}_p\left[(s-\mu_g)^2\right] &= (\mu_p - \mu_g)^2 + \sigma_p^2 \\
\mathbb{E}_p\left[(s-\mu_p)^2\right] &= (\mu_p - \mu_p)^2 + \sigma_p^2 = \sigma_p^2
\end{aligned} \tag{57.29}$$

Substituting expectations:

$$KL\big(P(s)\,\|\,G(s)\big) = \log_2\left(\frac{\sigma_g}{\sigma_p}\right) + \frac{(\mu_p-\mu_g)^2 + \sigma_p^2}{2\ln(2)\sigma_g^2} - \frac{\sigma_p^2}{2\ln(2)\sigma_p^2} \tag{57.30}$$

To describe error, we include an empirical variable β in this equation that relates the KL divergence (where $\beta = 1$) to an observed empirical error rate:

$$KL\big(P(s)\,\|\,G(s)\big) = \beta \cdot \left(\log_2\left(\frac{\sigma_g}{\sigma_p}\right) + \frac{(\mu_p-\mu_g)^2 + \sigma_p^2}{2ln(2)\cdot\sigma_g^2} - \frac{1}{2ln(2)}\right) \tag{57.31}$$

Here, μ_g and σ_g signify the mean and standard deviation of the goal distribution. μ_p and σ_p signify the mean and standard deviation of the progress distribution. The natural logarithms are constants and are introduced by base conversion. Including these allows for an exact result in bits of relative entropy. The main difference between Equations 57.15 and 57.31 is the addition of the log variability ratio (Senior et al., 2020) of goal and progress distributions. As such, Equation 57.31 models the *mismatch* in variances as an increase in the KL divergence that causes the average entropy to increase, raising the entire loss function.

57.5.1 Precision/Speed Trade-off

Variance σ_g^2 represents the base noise scale of the scaled goal tolerance *(zN)*, while variance σ_p^2 describes the uncertainty of the progress distribution towards that goal. Together, these quantities capture a precision/speed trade-off. The relationship to Fitts' law is evident: when movement speed exceeds what the goal tolerances permit, the movement becomes suboptimal, either by increasing endpoint error or by requiring corrective time. Stated more generally, when the user moves outside the tolerated region of the task outcome, the task is executed sub-optimally.

57.6 EMPIRICAL STUDY

We now turn to a case study, where we empirically examine whether our general model can predict human performance in a particular situation. If it can, it provides evidence that the general model is worthy of further exploration. In particular, we want to verify, via an empirical study, whether the above postulates could successfully model a task that does not yet have a known model, and function as a real-time measurement tool. We used our simplified KL divergence to measure the information processing capacity of participants in a simulated target *avoidance* task: maintaining distance while driving behind a car. Our attempt to develop a better model of a driver's capacity is more than an academic exercise, as such a model can help modern car interfaces better recognise and mitigate problems of driver distraction (Groeger, 2000), by directly modeling cognitive load.

57.6.1 Study Literature Review

In-depth analysis of crash data shows that nearly one-third of all fatal and injury crashes involve driver inattention (Wundersitz, 2019). Lee et al. (2001) identified a 30% (310 ms) increase in reaction time when a speech-based email system is used, with implications for safety. Strayer et al. (2011) demonstrated that inattentional blindness

can be caused by concurrent cell phone usage while driving, significantly increasing the risk of an accident. In a study involving adult drivers, Strayer and Drews (2004) observed an 18% slower reaction time while conversing on a cell phone. Systematic reviews of distracted driving, such as those by Simmons et al. (2016, 2017) and Caird et al. (2014) overwhelmingly demonstrate that tasks that require the drivers' visual attention, such as dialing, texting, and in-car infotainment touchscreen use generate more risk than tasks that do not require visual attention, as these direct the driver's attentive resources away from the primary task (Niu et al., 2019). Researchers have shown particular interest in the car-following task due to its prevalence in driving. Wolfe et al.'s (2022) *information acquisition theory* modelled the information a driver acquires through peripheral vision and eye movements. Boer pioneered the use of a behavioural entropy measure as an index of workload, and suggested that information entropy (the average amount of information) is the basis for driving behaviour (Boer, 2000). Boer et al. (2005) used the sample entropy algorithm to gauge performance of drivers, as demonstrated through a driving simulator study. Senders et al. (1967) did pioneering work on *occlusion distance*: the distance a driver feels comfortable driving with their eyes shut. Kujala et al. (2016) developed this metric further as a way to measure visual demand while driving. Building on Wilde (1982)'s *risk homeostasis theory*, Lu et al. (2013) developed a *desired safety margin* model as a quantified index of risk perception during car-following. That model is based on the leading and following car speeds, and the relative distance between them. Zhu et al. (2016) analysed the driver's choice of headway (the time interval between two cars at a given speed) in various traffic conditions using a large database. Wu et al. (2021) proposed a risk repulsion factor F for car-following as determined by the time headway, relative speed, and safe spatial headway. The driver's cognitive load has also been assessed and captured as a metric during car-following (Summala, 2000). This is important especially when secondary tasks consume some or all of that load. Subjective questionnaires are often carried out to collect the participants' assessment of task difficulty, risk, comfort, and effort. For example, Lewis-Evans et al. (2010) identified a negative relationship between the time headway in a car-following task and ratings of task difficulty in a simulator experiment. However, there are many issues with subjective assessments, including but not limited to bias caused by participants' preferences towards socially desirable responses (Lajunen & Summala, 2003). Others have proposed various physiological measurements as indicators of the subject's mental effort, such as heart rate variability, electromyogram, and facial temperature (de Waard, 1996; Kajiwara, 2014; Mulder, 1992). Yet such portable physiological sensing solutions are unreliable, while trajectory-based risk assessment methods are too complex (Petelczyc et al., 2020).

57.6.2 KL Divergence of Overlap

We used our framework to model the car-following task probabilistically, as illustrated in Figure 57.3. The task is represented using two normal distributions that describe the uncertainty in the lead and following car's position. In our model, these are the goal distribution (of the front car) and the progress distribution (of the following car). The task for the driver of the following car is to ensure the progress distribution avoids overlap with the goal distribution. The overlap between the two distributions follows a KL divergence that indicates the amount of bits of processing by the following driver to perform this task. Our simplified KL divergence models this as follows:

$$H = \beta \cdot \left(\frac{\sigma_g}{\mu_p - \mu_g} \right)^2 = \beta \cdot \left(\frac{N}{S} \right)^2 \qquad (57.32)$$

where μ_g is the mean of the goal distribution and μ_p is the progress distribution. Recognise that what matters here is the mean relative distance between cars $(\mu_p - \mu_g)$ as it relates to variability in the goal, the deceleration behaviour of the front car σ_g that needs matching. Note also that since we are

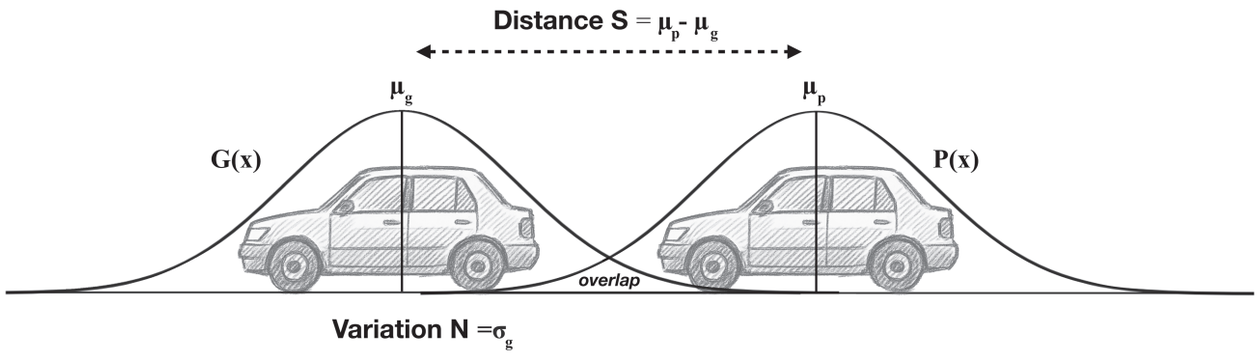


FIGURE 57.3 Two cars with normal distributions $G(x)$ and $P(x)$ representing uncertainty in their position. The percentage of overlap between distributions equals the probability of a collision. The negative logarithm of this probability gives the number of bits of difference.

avoiding rather than meeting the goal, we need to take the reciprocal of the SNR, i.e., the noise becomes the signal.

H is the amount of Bayesian surprise, in bits, that the driver would need to process in a car-following task. However, recall that the capacity C of the driver's ability to process this information is only logarithmic. We can establish parameters a and b of this logarithm by measuring the amount of information gained H using Equation 57.32 and by performing a regression on the result:

$$C = a + b \cdot log_2\left(\frac{z\sigma_g}{\mu_p - \mu_g}\right) = a + b \cdot log_2\left(\frac{zN}{S}\right) \quad (57.33)$$

57.6.3 Task

We designed a controlled car-following task in a driving simulator to collect driving data in target avoidance conditions with varying z-scaled noise-to-signal ratios. Our goal with this experiment was to model the drivers' information load using our equations to gauge the amount of information used to maintain a distance S when following a lead car with variation z times N (see Figure 57.3). We focused our experimental analysis on deceleration events that may lead to collision.

S and zN were varied between conditions. As in real driving, zN was not static within a condition: the lead car accelerated and decelerated to different speeds, which added a normally distributed noise zN to each distance condition, at different levels of z. This allowed us to capture participants' handling of the car in a range of driving difficulty levels defined by the constituting N/S ratios. We hypothesised that driver's information capacity during decelerations would be logarithmic with the noise-to-signal ratio, with a mean capacity b, in bits, of information per bit of difficulty. Because driver information capacity was logarithmic, but the amount of information generated was quadratic, at capacity b, the amount of information would outstrip the driver's ability to process. We hypothesised that this would lead to an error statistic that follows the KL divergence in Equation 57.32.

57.6.4 Participants

We recruited 15 participants (5 female, 10 male) from our organisation, with a mean participant age of 27. Four of the participants had low driving experience (one year or less).

57.6.5 Experimental Setup

We presented a simulated driving task with simulated road conditions using BeamNG (BeamNG, 2013), a car simulator video game chosen for its realism, physics engine, low cost, and programmability. Each participant driver was seated in a simulator in front of a curved 49-inch monitor with a 32:9 screen ratio. The monitor was positioned 0.6 m from the participant at eye level to provide an immersive simulator view of a car dashboard and windscreen (see Figure 57.4). A



FIGURE 57.4 Participant's experimental setup. The white bar in the centre of the screen helped participants perceive the correct distance S. They needed to match the size of this bar with the lead car's bumper while driving. Here, it is wider than the lead car's bumper indicating that the participant needs to accelerate.

Logitech G29 steering wheel and pedal unit with three pedals were mounted under the simulator rig. Participants were asked to adjust their seat position so that the steering wheel and accelerator/brake pedals were at a comfortable distance. The pedal clutch was not used. Every participant used the same custom car model with the same physics characteristics in BeamNG.

We constituted our conditions through custom scripts in BeamNG. We also wrote a Python program to interface with BeamNG to administer the study and to capture driving data at a 450 Hz sampling rate: vehicle positions and speeds of the driver's car and lead car, pedal data, and collision data. Our custom Python scripts filtered and analysed the data.

During the experiment, participants had a clear view of the lead car and road through their simulated windscreen. Because a 2D simulator did not provide 3D depth cues, participants were shown a floating fixed-width rectangular bar in the centre of their windscreen, positioned below the bumper of the lead car. That bar's width reflected what the perceived width of the lead car's bumper should be at the optimal following distance.

The driver's task was to maintain an optimal following distance, where the driver would try to match the bar's size to the perceived size of the bumper of the lead car. For example, in Figure 57.4, the bar is wider than the bumper, meaning the driver should move closer to the lead car.

57.6.6 Experiment Design

We used a within-subject experiment design where the task was to follow the lead car at a specified distance *S*. We then varied the speed of the lead car with variation *zN*. Each participant experienced 12 different task conditions with two factors *S* and *zN*.

- Factor *S* comprised four levels of fixed distances to the lead car (2.84 m, 4.84 m, 6.84 m, and 8.84 m).
- Factor *zN* comprised three levels centred around a mean lead car speed of 100 km/h. Each level represented a 4.1σ of a normal distribution of varying distances to the lead car, as generated by increasing the range of the lead car's allowable speeds (1.93 m: 100±12 km/h; 4.19 m: 100±18 km/h; and 7.21 m: 100±24 km/h). That is, the levels represented by factor *z* created increasing variability and thus noise in how the car in front accelerated and decelerated. To clarify these numbers, at level 1, the speed was allowed to vary between 88 km/h and 112 km/h (i.e., 100±12 km/h), where 96% of the variation in distance was captured by 1.93 m.

To constitute each distribution, for each level of *zN*, the lead car randomly decelerated 20 times and randomly accelerated 20 times, although always keeping within the level's range of speed values. Each trial within a level of *zN* comprised a new speed adjustment, which occurred approximately every two seconds. We randomised the presentation order of conditions. We also randomised the presentation order of acceleration or deceleration trials, with accelerations never occurring within a deceleration trial or vice versa.

We measured the instantaneous noise-to-signal ratio (change in distance by lead car ΔS, i.e., some sample of *N*; mean distance to lead car $\bar{S}$, per trial) and pedal activity by the driver. As will be detailed further, our dependent variables were the amount of information processed during a trial (KL divergence) and error (collisions). Together, the 40 trials created a condition simulating a continuous driving task of 2 minutes.

57.6.7 Procedure

Participants were instructed to hold but not use the steering wheel. They were instructed to only use the brake and accelerator pedals to maintain distance *S*, as the lead car varied its speed, by matching the bar with the bumper of the lead car. The road simulation was straight and without obstacles, allowing cars to travel in a straight line. Participants were given a 5-minute training task to familiarise themselves with the controls and task. Each condition began with both cars automatically accelerating to 100 km/h, with the specified distance *S*. Participants were then given a visual alert to take control of the car. After each condition ended, we asked the participants to rate the driving difficulty of that condition using a 7-point Likert scale (Lewis-Evans et al., 2010). Trials that resulted in collision were repeated, with any collisions only counting once per trial, per participant.

57.6.8 Ecological Validity

Our task was designed such that participants would be in a position to experience error. A choice for longer distances and lower speeds would mean the chances of a collision would be too low to be measurable and compared with the driver's remaining capacity. This meant having relatively low headways at high absolute speeds that could be criticised for their limited ecological validity. We note, however, that the important parameter was not the absolute speed at which the cars were driving, but the variation in their average relative distance over time. This applies equally to parking a car with respect to a solid barrier at low speeds, as it does to avoiding a lead car at 100 km/h. We also chose for the simulator rendering to have most environmental cues removed, making it difficult for participants to estimate an absolute speed. This choice made the experiment more generalisable by avoiding potential confounding factors generated by environmental renderings. The choice not to have participants steer was inspired by the potential for introducing confounding variables as well.

57.7 RESULTS

We first describe our data analysis, and then present the results of our empirical evaluation of the models.

57.7.1 Measurement and Data Analysis

In our experiment, lead car accelerations never resulted in a collision. Given our interest in the relationship between target avoidance and the risk of collision (error), we constrained our analysis to trials where participants responded to a deceleration of the lead car by braking or by releasing the accelerator pedal. We used three different levels of tolerance for noise-to-signal measurements in our analysis. Each used a different confidence interval z for defining zN.

1. At the condition level of precision, zN was defined by 4.1σ of the distribution of relative backwards distances travelled by the lead car during trials in that condition;
2. At the trial level zN was defined by the actual relative backwards distance travelled by the lead car during each trial; and
3. Within each trial, N was defined by the standard deviation (σ) of a one-second time window of instantaneous distance measurements, as detailed below.

Empirical Factors Noise (zN) and Signal (S). Although we used predetermined levels of factor S, we administered factor zN by altering the speed of the lead vehicle, which was affected by the physics engine. As such, we measured the actual negative relative distance travelled by the lead car (ΔS) per trial, from the moment of deceleration of the lead car to the lead car reaching the desired speed. In line with Welford's (1968) statistical analysis of Fitts' law target width W, we then calculated zN as 4.1σ of the distribution of the ΔSs actually administered per level of condition zN.

Empirical Noise-to-Signal Ratio. As each trial involved a braking event that resulted in processing of information given an individually administered ΔS, many similar zN/S ratios were created across conditions. We used these more detailed empirical observations of noise-to-signal ratios for our log-linear regression analysis, pooled across all trials. Here, our definition of noise-to-signal ratio was $N = \Delta S$ per trial divided by the measured mean S ($\bar{S}$) per trial, yielding the metric $\Delta S/\bar{S}$. We also used this definition for our regression analysis of error rates, as this yielded directly comparable results of error rates vs. information gain in bits. This definition was also used for our index of difficulty (see Equation 57.36). Note that the only effect of the use of sampled ΔS values instead of some confidence interval N is a scaling effect on zN/S ratios.

Measuring Information Processed. Our first dependent variable was the amount of information processed by the driver during decelerations. To measure the instantaneous amount of information processed within a trial, we used a more detailed one-second rolling time window to sample the distribution of distances S to the lead car at 450 Hz. Our base definition of noise N here was the standard deviation (σ) of the distribution of S obtained in this time window, while our definition of signal S was the mean S ($\bar{S}$) of this time window. First, the noise-to-signal ratio was sampled using this type of one-second window at the beginning of each trial, ending at the moment the driver responded to the stimulus. We then calculated the amount of information in the stimulus by entering this first noise-to signal ratio into Equation 57.32. A one-second window again sampled the noise-to-signal ratio at the end of each trial, prior to (i.e., ending at) a new deceleration or acceleration event. We calculated the remaining amount of information by entering this second noise-to-signal ratio into Equation 57.32.

Subtracting the remaining amount of information from the amount of information in the stimulus gave us the amount of information processed by the driver during each deceleration trial. We subsequently aggregated measurements across all conditions, sorting results using $\Delta S/\bar{S}$, and binning measurements that were within 0.05 units of $\Delta S/\bar{S}$.

Measuring Error. Our second dependent variable was error. Error was defined, across conditions, as the occurrence of a collision with the lead car. All trials that ended in error were repeated. Since any instantaneous measure of N/S during a collision would yield infinity during the one-second period leading up to collision, we used the $\Delta S/\bar{S}$ of the repeat of the trial as our signal-to-noise metric. This allowed comparison of error and capacity metrics using the same scale. We could not establish any measure for unprocessed information for trials that ended in collision. We could, however, convert error rates into bits of information in a manner that provided insights into the relationship between driver capacity and error rates. After binning errors within 0.5 $\Delta S/\bar{S}$, we calculated the error probability $p(E)$ by dividing the number of errors per bin by the total number of errors across all conditions. We then converted the complement of this probability to bits E_b (the non-surprise of a collision) using a negative log probability transform (Shannon, 1948):

$$E_b = -\log_2\big(1 - p(E)\big) \tag{57.34}$$

57.7.2 Effects of Factors zN and S

Table 57.1 shows the mean KL divergence values for consumed information against levels of factors zN and S. Since the residuals in our ANOVA model of the dependent variable differed significantly from a normal distribution (KS test, $p > 0.05$), we used a nonparametric Friedman test. The Friedman test showed highly significant effects of both factors zN ($\chi^2(2) = 17.73$, $p < 0.001$) and S ($\chi^2(3) = 40.84$, $p < 0.001$) on the number of bits consumed by the driver (note that a two-way ANOVA produced the same results).

TABLE 57.1
Mean KL (s.e.) for Consumed Surprise (in Bits) for Levels of Factors *zN* and *S*

N	2.84 m	4.84 m	6.84 m	8.84 m
1.93 m	0.11 (0.07)	0.06 (0.04)	0.03 (0.02)	0.01 (0.02)
4.19 m	0.16 (0.09)	0.09 (0.04)	0.05 (0.03)	0.04 (0.04)
7.20 m	0.17 (0.09)	0.10 (0.03)	0.08 (0.04)	0.05 (0.02)

57.7.3 Car-Following Model

Log-linear regression of binned ΔKL divergence measurements onto noise-to-signal measurements $\Delta S/\bar{S}$ yielded a logarithmic relationship ($r^2 = 0.86$, F (1,26) = 164.2, $p < 0.001$):

$$C = -0.02 + 0.25 \log_2 \left(\frac{\Delta S}{\bar{S}} + 1 \right) \tag{57.35}$$

Figure 57.5 shows this relationship between driver capacity (information gain) and aggregated signal-to-noise measurements on a $\Delta S/\bar{S}$ scale across conditions, as well as on a logarithmic scale (below the *x*-axis). Mean capacity is visualised as a dashed line. Subtracting driver capacity from this mean capacity yields remaining driver capacity (black solid line) i.e., remaining mental load. The logarithmic scale at the bottom of Figure 57.5 shows the index of difficulty (ID) of this task in bits:

$$ID = \log_2 \left(\frac{\Delta S}{\bar{S}} + 1 \right) \tag{57.36}$$

The mean capacity (bandwidth *b*) of our car-following task was measured at $b = 0.25$ bits/bit, while the index of performance (throughput 1/*b*) was 3.96 bits/bit. The *y*-axis indicates bits of consumed information, while the *x*-axis indicates noise-to-signal ratio $\Delta S/\bar{S}$, or bits of *ID* if the bottom log scale is used. Note that the only difference in plotting Figure 57.5 on a log-linear scale against *ID*, is that lines would be straight, with the bottom scale as the primary *x*-axis. Mean capacity is expressed as bits of surprises consumed per bit of difficulty, while the index of performance (1/*b*) is expressed as the bits of difficulty processed per bit of surprise.

57.7.4 Error (Collision) Model

Regression of error metric E_b (non-surprise of a collision in bits) onto noise-to-signal measurements ($\Delta S/\bar{S}$) yielded a significant relationship with the KL divergence metric (Equation 57.32) ($r^2 = 0.94$, F (1,5) = 80.5, $p < 0.001$):

$$E_b = 0.03 + 0.13 \left(\frac{\Delta S}{S} \right)^2 \tag{57.37}$$

Figure 57.6 shows the relationship between error (large dash) and aggregated signal-to-noise measurements on a $\Delta S/\bar{S}$ scale across conditions, as well as on a logarithmic scale (*ID*, bottom), compared with driver capacity (dotted line) and remaining capacity (black solid). The total number of collisions across all trials was 90. Error rates below or equal to a noise-to-signal ratio or *ID* of 1 trended, on

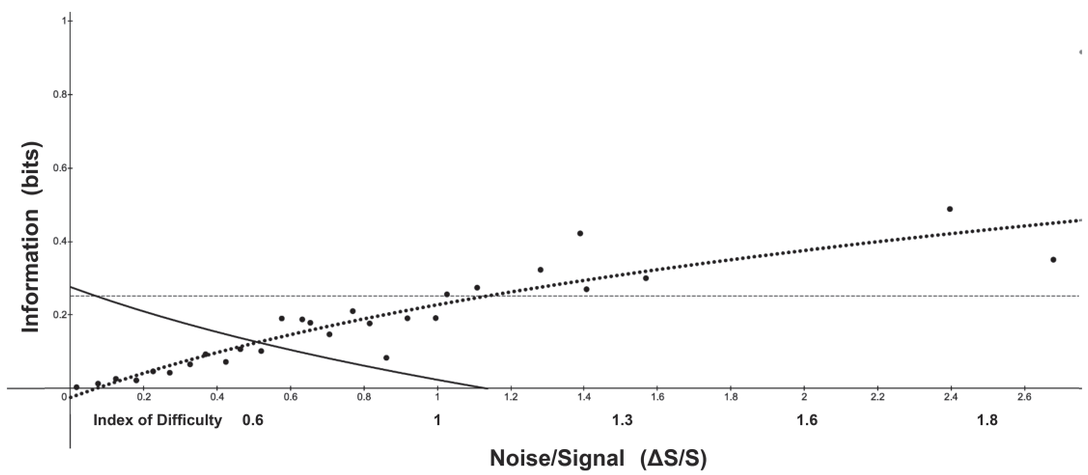


FIGURE 57.5 Driver information capacity plotted against noise-to-signal ratio ($\Delta S/S$, *ID* on a logarithmic scale below *x*-axis). Capacity followed a logarithmic function (dotted line), with a slope of 0.25 bits. This average capacity is also drawn as a small dashed line. Black solid line indicates remaining capacity.

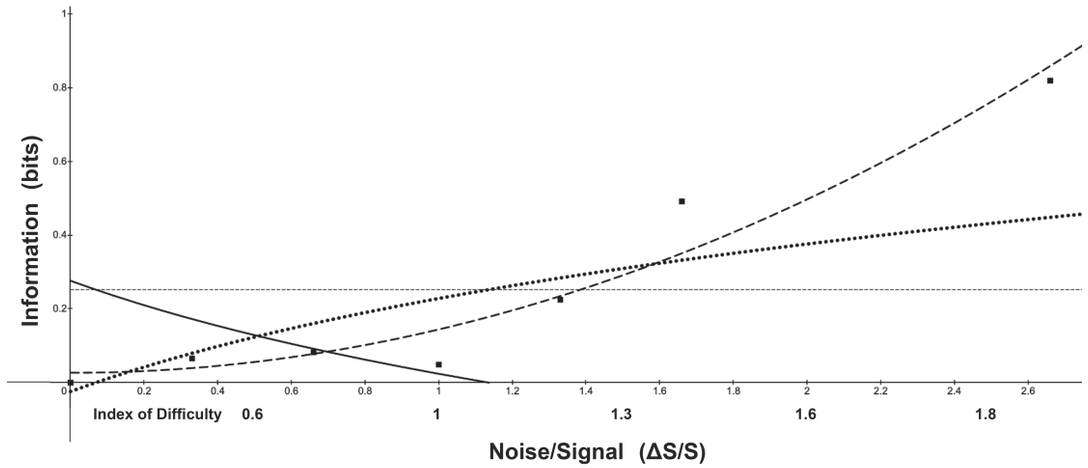


FIGURE 57.6 Error (non-surprise in bits) plotted against noise-to-signal ratio (Δ*S/S*, *ID* on a logarithmic scale below *x*-axis). Error followed a KL divergence (large dash). Dotted line shows driver capacity and the black solid line represents remaining capacity. The small dash shows average capacity.

average, slightly below the expected 4% (at 0.05 bits or 3.4%, between 2 and 4 collisions across trials), while those above 1 increased quadratically, averaging 0.38 bits (23%), with a maximum of 0.82 bits (43% or 39 collisions across trials) at *ID* = 1.85.

57.7.5 Qualitative Results

Figure 57.7 shows participants' questionnaire responses ranking difficulty per conditional *zN/S*, sorted by *ID*. These represent mean responses to the question "This driving task was difficult", on a scale from 1 (Strongly Disagree) to 7 (Strongly Agree) (Lewis-Evans et al., 2010). A Friedman non-parametric analysis of variance showed participants rated the car-following task as significantly more difficult with increasing *ID* ($\chi^2(11) = 116.4$, $p < 0.0001$). Note that all participants found the task to be easy to interpret and execute, while none reported distraction by the floating bar graphic during the task.

57.8 DISCUSSION

We will now discuss the results from our empirical study, relate them to our interactive inference model, and provide some suggestions for future directions.

57.8.1 Empirical Study

The above results are a first step towards empirical validation of both our capacity and error models. Results show a good-to-excellent fit between predicted and observed models of the task. Driver information processing capacity during deceleration events in a simulated car-following task can be measured in real time using the simplified KL

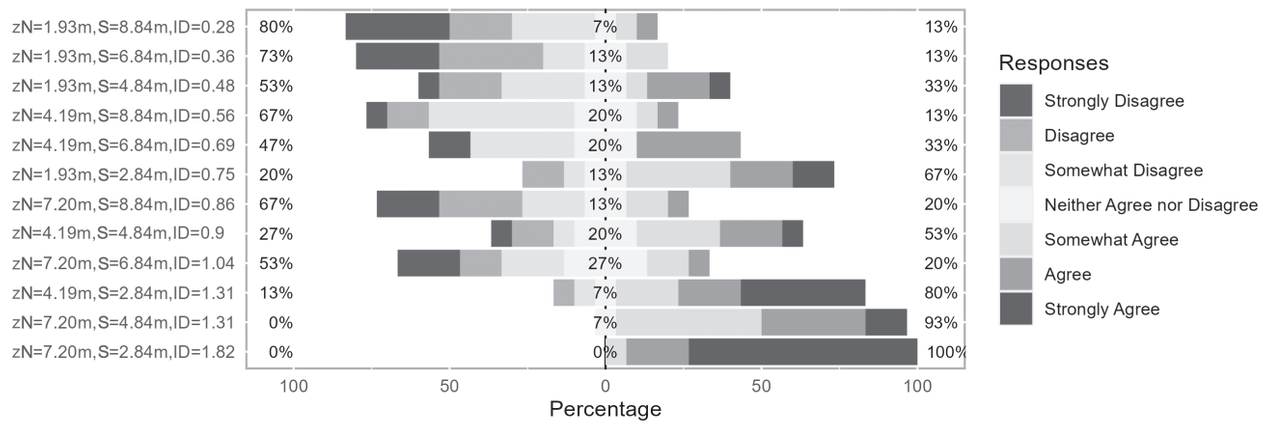


FIGURE 57.7 Perceived task difficulty for each condition: Participants' responses to the question "This driving task was difficult" on a scale of 1 (Strongly Disagree) to 7 (Strongly Agree).

divergence of the noise-to-signal ratio of the distance to the car followed, over a time window. Our hypotheses that both distance *S* and its variance *zN* to the lead car significantly affect the amount of information consumed by the driver during such events were confirmed, with more information consumed at lower *S*s and higher *zN*s.

Error rates also appear to independently confirm the effect of *zN/S* ratios on driver information gain during car-following. As expected, the (negative log of the complement of) error rate followed a KL divergence. When driver capacity exceeds mean capacity, error rates appear to increase with the square of the noise-to-signal ratio. We attribute these trends to "left-over" Bayesian surprise, generated by high noise-to-signal ratios, that is not consumed by the driver. Note that we used mean capacity as our capacity limit because the theoretical maximum channel capacity is, in practice, never reached (without significant error) (Shannon, 1948). Our results also suggest bit rates are quite low. We presume this is because information processing is throttled by the eye/foot/pedal processing loop.

Qualitative ratings by participants were in line with the above. They rated the task as more difficult with higher noise-to-signal ratios (and thus *ID*s), albeit with bipolar ratings of middle *ID*s. We attribute this to participants perceiving distant targets ($S > 4.8$ m) as easier to process than close targets with similar SNRs, perhaps due to the 2D nature of our simulator setup. We again note that our results are limited to modelling driver responses to lead car decelerations in a simulated car-following task using specific road conditions, car physics, and a sample set of participants.

57.8.2 Interactive Inference Model and Future Directions

Our results provide good support for the interactive inference model, which shows how capacity for performance can be related to error, both in bits. We model capacity (information gain) as a logarithmic function of SNR, scaled by a slope that is empirically determined. We can consider this slope to be analogous to the bandwidth in the Shannon-Hartley theorem. Because capacity beyond the limits of this model cannot occur without error and because both metrics can be converted to bits, one of the predictions from our model is that error can be analysed in view of left-over capacity. This is because the average capacity is, in fact, the maximum capacity without error. Error subsequently tracks Bayesian surprise as a KL divergence. Figure 57.6 shows this relationship, with the KL divergence fitting an empirically observed error rate (converted to negative log probability) visualised with a dashed line. Capacity is modelled as a logarithmic function (dotted line), which is an empirical fit to Bayesian surprise processed at different SNRs of the task. By subtracting the logarithmic function from this value, we can arrive at a left-over capacity metric (in black). Where this line reaches zero capacity, errors increase and start tracking the KL divergence: the square of the SNR. The empirical evaluation of this relationship is considered a future direction.

Left-over capacity would correspond directly to the cognitive or sensorimotor load of the task. By correlating empirical results with heart rate variability (HRV), data and other measures of cognitive load evidence for this relationship could be examined. One important topic of investigation is to determine if cognitive load could be measured during tasks in real time. For example, users could complete a task modelled via our method while wearing a heart rate monitor. HRV should be lower during moments in the task where there is significant left-over surprise.

The correlation between HRV and cognitive load suggests it may be possible that leftover Bayesian surprise also tracks stress levels in participants. The relationship between our models and experienced stress levels of participants during tasks is therefore an interesting area of future study. For example, users could complete a task modelled via our method and then fill out a questionnaire on stress levels. Left-over surprise metrics should correlate with questionnaire results if our hypothesis is correct.

Because our metrics are in bits, rather than time, it should be possible to compare both capacity and error between different tasks. Analysing different tasks for an SNR parameter that can be measured, and comparing results from empirical investigations of the relationship between SNR and both capacity and error, would further evidence this assertion. For example, users could complete different tasks while wearing HRV monitors. Left-over capacity metrics should correlate not just with HRV data, but also between tasks, if our hypothesis is correct. We consider such empirical work a future direction and invite others to help investigate this promising area of research.

57.9 LIMITATIONS

This chapter presents an exploratory framework inspired by active inference principles, which does not require implementation of its full mathematical machinery (such as model evidence and expected free energy). Our use of the term Bayesian surprise refers specifically to the KL divergence between our operationally defined goal and progress distributions. In addition, the connections we draw to Fitts' law, Hick's law, and the power law are observational. Finally, our empirical validation is limited to a single driving task, and further research is needed to test the generalisation of this approach to other interaction contexts.

57.10 CONCLUSIONS

Neuromorphic HCI aims to improve user interfaces by exploring how concepts from neuroscience might inform design. Interactive inference offers a simplified interpretation of the neuroscience theory of active inference in user interface design. It postulates that users perform Bayesian inference on current progress and goal distributions to learn tasks and predict their next action. The Bayesian surprise (relative entropy) between goal and progress distributions can be modelled as a square function. However, human

capacity to process Bayesian surprise follows a logarithmic function. This explains why errors occur when the user runs out of capacity. We believe our model may allow the quantitative analysis of performance and error in a single framework that allows estimation of the mental load of a task. We showed how three basic laws of HCI – Hick's law, Fitts' law, and the power law of practice – can be derived from interactive inference. We evaluated the use of the model in an empirical study. In our simulated car-following task, driver information processing capacity during braking events was linearly related, with good fit, to a logarithmic model of the noise-to-signal ratio of the distance to a lead car, plus 1 (in bits). We found a significant effect of both distance S and tolerance zN on capacity, with a mean driver index of performance of almost four bits of difficulty (as defined by $log_2(zN/S+1)$) per bit of surprise. Participants' qualitative ratings followed the same trend, with significantly higher ratings of difficulty with higher zN/S ratios. We used our model to derive a remaining capacity for the driver, by subtracting the log-linear curve from mean capacity, in bits. We also fit error (collision) rates, with excellent fit, to the square of the signal-to-noise ratio of the distance to the lead car in bits. We hypothesise that collisions are due to unprocessed surpirse where driver capacity is exceeded. Overall, the empirical results provide support for our interactive interference model of HCI. While car-following is just one case study, it provides a framework for conducting further empirical research in interactive inference, and evidence that such research is well worth doing. We invite other researchers to test and apply this model to other tasks, to see how well these models predict human performance, and to confirm or refute the general applicability of neuromorphic HCI in general, and interactive inference in particular.

REFERENCES

Anderson, J. R., Matessa, M., & Lebiere, C. (1997). ACT-R: A theory of higher-level cognition and its relation to visual attention. *Human–Computer Interaction*, *12*(4), 439–462.

Baldi, P., & Itti, L. (2010). Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Networks*, *23*(5), 649–666.

Bayes, T. (1763). An essay towards solving a problem in the doctrine of chances. *Philosophical Transactions of the Royal Society*, *53*, 370–418.

Beal, M. J. (2003). *Variational algorithms for approximate Bayesian inference*. Doctoral dissertation, University of London, London.

BeamNG. (2013). *Beamng.drive*. https://www.beamng.com/game/

Bishop, C. M., & Bishop, H. (2023). *Deep learning: Foundations and concepts*. Springer Nature.

Boer, E. R. (2000). Behavioral entropy as an index of workload. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, *44*(17), 125–128.

Boer, E. R., Rakauskas, M. E., Ward, N. J., & Goodrich, M. A. (2005). Steering entropy revisited. In *Proceedings of the third international driving symposium on human factors in driver assessment, training and vehicle design*. (Vol. 3, pp. 25–32).

Boltzmann, L. (1866). *Über die mechanische bedeutung des zweiten hauptsatzes der wärmetheorie*. Staatsdruckerei.

Caird, J., Johnston, K., Willness, C., Asbridge, M., & Steel, P. (2014). A meta-analysis of the effects of texting on driving. *Accident Analysis & Prevention*, *71*, 311–318.

Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human-computer interaction*. L. Erlbaum Associates.

Clausius, R. (1879). *The mechanical theory of heat*. Macmillan.

Cohen, J. (1988). *Statistical power analysis for the behavioral sciences* (2nd ed.). Lawrence Erlbaum Associates.

Courant, R. (1943). Variational methods for the solution of problems of equilibrium and vibrations. *Bulletin of the American Mathematical Society*, *49*(1), 1–23.

Crossman, E. R. F. W. (1959). A theory of the acquisition of speed-skill. *Ergonomics*, *2*(2), 153–166.

de la Fuente, H. L. (2014). *Etude multi-niveaux du contrôle d'un périphérique d'interaction de type joystick*. Doctoral dissertation, Aix-Marseille Université, France.

de Waard, D. (1996). *The measurement of drivers' mental workload*. Doctoral dissertation, University of Groningen, The Netherlands.

Feynman, R. P. (1972). *Statistical mechanics: A set of lectures*. W. A. Benjamin.

Fitts, P. M. (1954). The information capacity of the human motor system in controlling the amplitude of movement. *Journal of Experimental Psychology*, *74*, 381–391.

Friston, K. (2012). A free energy principle for biological systems. *Entropy*, *14*(11), 2100–2121.

Friston, K., Kilner, J., & Harrison, L. (2006). A free energy principle for the brain. *Journal of Physiology-Paris*, *100*(1–3), 70–87.

Groeger, J. A. (2000). *Understanding driving: Applying cognitive psychology to a complex everyday task* (1st ed.). Routledge.

Guiard, Y., Olafsdottir, H. B., & Perrault, S. T. (2011). Fitts' law as an explicit time/error trade-off. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 1619–1628). North Holland.

Hart, S. G., & Staveland, L. E. (1988). Development of NASA-TLX: Results of empirical and theoretical research. In P. A. Hancock & N. Meshkati (Eds.), *Human mental workload* (pp. 139–183). Elsevier North-Holland.

Hick, W. E. (1952). On the rate of gain of information. *Quarterly Journal of Experimental Psychology*, *4*(1), 11–26.

Hyman, R. (1953). Stimulus information as a determinant of reaction time. *Journal of Experimental Psychology*, *45*(3), 188–96.

Kajiwara, S. (2014). Evaluation of driver's mental workload by facial temperature and electrodermal activity under simulated driving conditions. *International Journal of Automotive Technology*, *15*, 65–70.

Knuth, D. E. (1998). *The art of computer programming, volume 3: Sorting and searching* (2nd ed.). Addison-Wesley.

Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences*, *10*(7), 319–326.

Kujala, T., Mäkelä, J., Kotilainen, I., & Tokkonen, T. (2016). The attentional demand of automobile driving revisited: Occlusion distance as a function of task-relevant event density in realistic driving scenarios. *Human Factors*, *58*(1), 163–180.

Kullback, S. (1997). *Information theory and statistics*. Dover Publications.

Lajunen, T., & Summala, H. (2003). Can we trust self-reports of driving? Effects of impression management on driver behaviour questionnaire responses. *Transportation Research Part F: Traffic Psychology and Behaviour*, *6*(2), 97–107.

Lee, J. D., Caven, B., Haake, S., & Brown, T. L. (2001). Speech-based interaction with in-vehicle computers: The effect of speech-based e-mail on drivers' attention to the roadway. *Human Factors: The Journal of Human Factors and Ergonomics Society*, *43*(4), 631–640.

Lewis-Evans, B., De Waard, D., & Brookhuis, K. A. (2010). That's close enough—A threshold effect of time headway on the experience of risk, task difficulty, effort, and comfort. *Accident Analysis & Prevention*, *42*(6), 1926–1933.

Lu, G., Cheng, B., Wang, Y., & Lin, Q. (2013). A car-following model based on quantified homeostatic risk perception. *Mathematical Problems in Engineering*, *2013*, 1–13.

MacKenzie, I. S. (1992). Fitts' law as a research and design tool in human-computer interaction. *Human-Computer Interaction*, *7*(1), 91–139.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, *88*(5), 375.

Moulin, P., & Johnstone, P. R. (2014). Kullback-leibler divergence and the central limit theorem. In *UCSD information theory and applications workshop.*

Mulder, L. J. (1992). Measurement and analysis methods of heart rate and respiration for use in applied environments. *Biological Psychology*, *34*, 205–236.

Niu, J., Wang, X., Liu, X., Wang, D., Qin, H., & Zhang, Y. (2019). Effects of mobile phone use on driving performance in a multiresource workload scenario. *Traffic Injury Prevention*, *20*(1), 37–44.

Novak, R. (2015). Lecture 7: Hypothesis testing and KL divergence. *Lecture Notes.*

Petelczyc, M., Gieraltowski, J., Zogal a Siudem, B., & Siudem, G. (2020). Impact of observational error on heart rate variability analysis. *Heliyon*, *6*(5).

Prigogine, I. (1977). *Self-organization in nonequilibrium systems: From dissipative structures to order through fluctuations*. Wiley.

Rao, R. P. N., & Ballard, D. H. (1999). Predictive coding in the visual cortex: A functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, *2*(1), 79–87.

Reason, J. (1990). *Human error.* Cambridge University Press.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning representations by backpropagating errors. *Nature*, *323*(6088), 533–536.

Schmidt, R. A., Zelaznik, H., Hawkins, B., Frank, J. S., & Quinn, J. T. Jr (1979). Motor-output variability: A theory for the accuracy of rapid motor acts. *Psychological Review*, *86*(5), 415.

Senders, J. W., Kristofferson, A. B., Levison, W. H., Dietrich, C. W., & Ward, J. L. (1967). The attentional demand of automobile driving. *Highway Research Record*, *195*(2), 15–33.

Senior, A. M., Viechtbauer, W., & Nakagawa, S. (2020). Revisiting and expanding the meta-analysis of variation: The log coefficient of variation ratio, lnCVR. *bioRxiv.* https://doi.org/10.1101/2020.01.06.896522

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, *27*(3), 379–423.

Simmons, S. M., Caird, J. K., & Steel, P. (2017). A meta-analysis of in-vehicle and nomadic voice recognition system interaction and driving performance. *Accident Analysis and Prevention*, *106*, 31–43.

Simmons, S. M., Hicks, A., & Caird, J. K. (2016). Safety-critical event risk associated with cell phone tasks as measured in naturalistic driving studies: A systematic review and meta-analysis. *Accident Analysis and Prevention*, *87*, 161–169.

Snoddy, G. S. (1926). Learning and stability: A psychophysiological analysis of A case of motor learning with clinical applications. *Journal of Applied Psychology*, *10*(1), 1.

Strayer, D. L., & Drews, F. (2004). Profiles in driver distraction: Effects of cell phone conversations on younger and older drivers. *Human Factors: The Journal of Human Factors and Ergonomics Society*, *46*, 640–649.

Strayer, D. L., Watson, J. M., & Drews, F. A. (2011). Chapter two – Cognitive distraction while multitasking in the automobile. In B. H. Ross (Ed.), *Advances in research and theory* (Vol. 54, pp. 29–58). Academic Press.

Summala, H. (2000). Brake reaction times and driver behavior analysis. *Transportation Human Factors*, *2*(3), 217–226.

Welford, A. T. (1968). *Fundamentals of skill*. Methuen Co Ltd.

Wilde, G. J. S. (1982). The theory of risk homeostasis: Implications for safety and health. *Risk Analysis*, *2*, 209–225.

Williamson, J. H., Oulasvirta, A., Kristensson, P. O., & Banovic, N. (2022). An introduction to Bayesian methods for interaction design. In *Bayesian methods for interaction and design* (pp. 3–80). Cambridge University Press.

Wobbrock, J. O., Cutrell, E., Harada, S., & MacKenzie, I. S. (2008). An error model for pointing based on Fitts' law. In *Proceedings of the SIGCHI conference on human factors in computing systems* (pp. 1613–1622).

Wolfe, B., Sawyer, B. D., & Rosenholtz, R. (2022). Toward a theory of visual information acquisition in driving. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, *64*(4), 694–713.

Wolpert, D. M. (2007). Probabilistic models in human sensorimotor control. *Human Movement Science*, *26*(4), 511–524.

Wu, B., Yan, Y., Ni, D., & Li, L. (2021). A longitudinal car-following risk assessment model based on risk field theory for autonomous vehicles. *International Journal of Transportation Science and Technology*, *10*(1), 60–68.

Wu, T., Dufford, A. J., Egan, L. J., Mackie, M.-A., Chen, C., Yuan, C., Chen, C., Li, X., Liu, X., Hof, P.R., & Fan, J. (2018). Hick–Hyman law is mediated by the cognitive control network in the brain. *Cerebral Cortex*, *28*(7), 2267–2282.

Wu, W., Gao, Y., Bienenstock, E., Donoghue, J. P., & Black, M. J. (2006). Bayesian population decoding of motor cortical activity using a Kalman filter. *Neural Computation*, *18*(1), 80–118.

Wundersitz, L. (2019). Driver distraction and inattention in fatal and injury crashes: Findings from in-depth road crash data. *Traffic Injury Prevention*, *20*(7), 696–701.

Zhu, M., Wang, X., & Wang, X. (2016). Car-following headways in different driving situations: A naturalistic driving study. In *International Conference of Transportation Professionals* (pp. 1419–1428). China: American Society of Civil Engineers.

BK-TandF-JACKO_9781032513485-260130-Chp57.indd 1426 27/06/26 10:26 AM