Recognizing Sights, Smells, and Sounds using Gnostic Fields

Christopher Kanan (ckanan@ucsd.edu)
Department of Computer Science and Engineering, University of California, San Diego
9500 Gilman Drive, La Jolla, CA 92093-0404

Mammals rely on vision, audition, and olfaction to remotely sense stimuli in their environment. Determining how the mammalian brain uses this sensory information to recognize objects has been one of the major goals of psychology and neuroscience. Likewise, researchers in computer vision, machine audition, and machine olfaction have endeavored to discover good algorithms for stimulus classification.

Almost 50 years ago, the neuroscientist Jerzy Konorski proposed gnostic fields as a theoretical model for object classification (Konorski, 1967). A gnostic field contains competing sets of gnostic neurons sitting atop sensory processing hierarchies, and Konorski argued that this would enable stimuli to be robustly categorized, despite variations in their presentation. Much of what Konorski hypothesized has been remarkably accurate (Gross, 2002), and neurons with gnostic-like properties have been discovered in visual (e.g., Desimone, 1991; Kreiman et al., 2000), aural (e.g., Averbeck & Romanski, 2006), and olfactory (e.g., Chapuis & Wilson, 2012) brain regions. I describe the first computational implementation of Konorski's gnostic field theory of object recognition across modalities (Kanan, 2013), and I assess its accuracy on large many-category datasets.

The model's ability to classify images is assessed on the challenging Caltech-256 and Caltech-UCSD Bird 200 (CUB-200) datasets, which both contain real-world images. The model was also used to classify musical artists using the Artist-20 dataset, and it was used to classify 108 chemicals sampled using an electronic nose. In all cases, gnostic fields exceeded state-of-the-art methods in computer vision, machine audition, and machine olfaction.

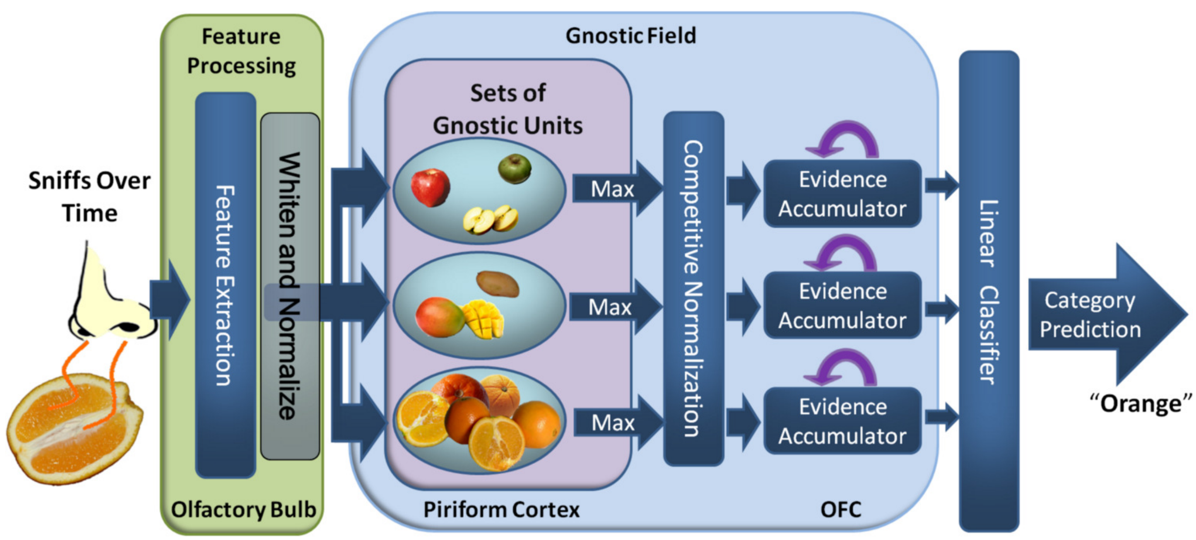

Figure 1. A cartoon depiction of an olfactory gnostic field for distinguishing among apples, mangos, and oranges. After whitening the olfactory features, units in gnostic sets compete to identify the input. This is followed by competitive normalization. The model accumulates evidence over time, which is combined using a linear classifier to determine what category is being observed.