

Artificial Intelligence (AI)

FPGA community forums and blogs have moved to the Altera Community. Existing Intel Community members can sign in with their current credentials.


Intel Community / Blogs / Tech Innovation / Artificial Intelligence (AI)

804 Discussions

Nature Machine Intelligence Publishes Intel Labs' Neuromorphic Research on Visual Perception

Subscribe

Article Options



Scott Bair Employee

07-15-2024

0 0 9,055

Posted on behalf of authors Paxton Frady and Friedrich T. Sommer, research scientists in the Neuromorphic Computing Lab at Intel Labs.

Highlights

- Two Intel Labs research papers present novel neuromorphic computing solutions using Intel's Loihi neuromorphic research chip to solve robotics visual perception problems with scene understanding and visual odometry in the June 2024 issue of Nature Machine Intelligence.
- Intel Labs collaborated with researchers at Forschungszentrum Jülich, University of Aachen, University of Zürich and ETH Zürich, ZHAW Wädenswil, Accenture Labs, and Redwood Center for Theoretical Neuroscience at UC Berkeley.
- Researchers used a new neural network architecture called resonator networks to efficiently solve inference problems through recurrent network dynamics. When combined with energy-efficient neuromorphic hardware, resonator networks can act as a perceptual engine for mobile devices. The research will enable drones, satellites, and other energy-constrained edge devices to analyze sensor input in ways that resemble perception in animals and humans.

Two Intel Labs research papers present novel neuromorphic computing solutions using Intel's Loihi neuromorphic research chip to solve robotics visual perception problems in scene understanding and visual odometry (VO). Published in the June 2024 issue of Nature Machine Intelligence, the studies resulted from collaborations between researchers at Intel Labs and Forschungszentrum Jülich, University of Aachen, University of Zürich and ETH Zürich, ZHAW Wädenswil, Accenture Labs, and Redwood Center for Theoretical Neuroscience at UC Berkeley.

The first paper develops a new approach to visual perception and scene understanding using resonator networks, a new type of neural network. The research describes how resonator networks can infer the identity and properties such as shape, color, and pose of objects in a scene from a given visual input. The companion paper demonstrates how resonator networks can be used in real-world tasks for visual odometry in robotics. Visual odometry uses data from cameras to estimate changes in position over time in robotics applications. When combined with energy-efficient Loihi neuromorphic hardware, resonator networks can act as a perceptual engine for mobile devices. The research will enable drones, satellites, and other energy-constrained edge devices to analyze sensor input in ways that resemble perception in animals and humans.

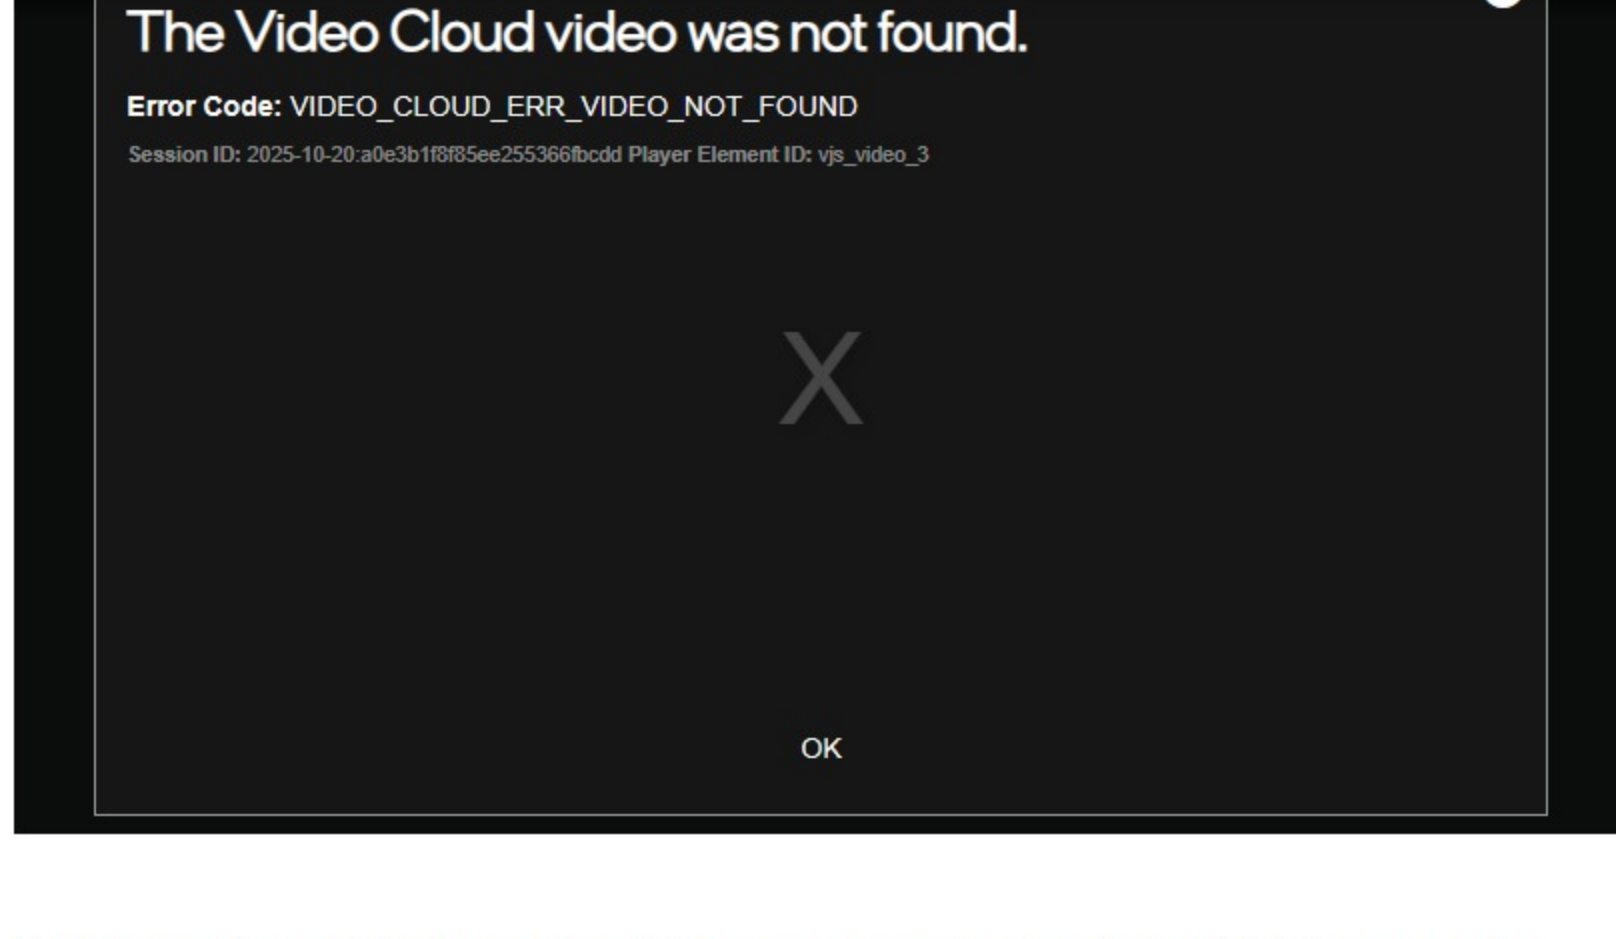
Conventional artificial intelligence (AI) methods typically require extensive training data using models with millions to billions of parameters, resulting in massive compute and energy costs while often remaining an unexplainable black box. Instead, researchers have proposed a transparent approach based on probabilistic inference in a generative model. To do inference efficiently, the research proposes a computational framework based on vector symbolic architectures (VSA). VSA is a transparent framework for parallel computing with distributed vector representations and is the foundation for creating efficient algorithms that benefit from neuromorphic hardware acceleration. Unlike most convention neural network approaches, VSAs include a binding operation, which allows for the expressivity required for the generative modeling approach. Using vector binding, finding transformations applied to objects in a scene or to the camera on a robotics platform can be cast as an inference problem. Resonator networks can solve such inference problems efficiently through recurrent network dynamics, requiring significantly smaller parameter counts than comparable deep learning networks.



(Image credit: Nature Machine Intelligence).

Visual Scene Understanding

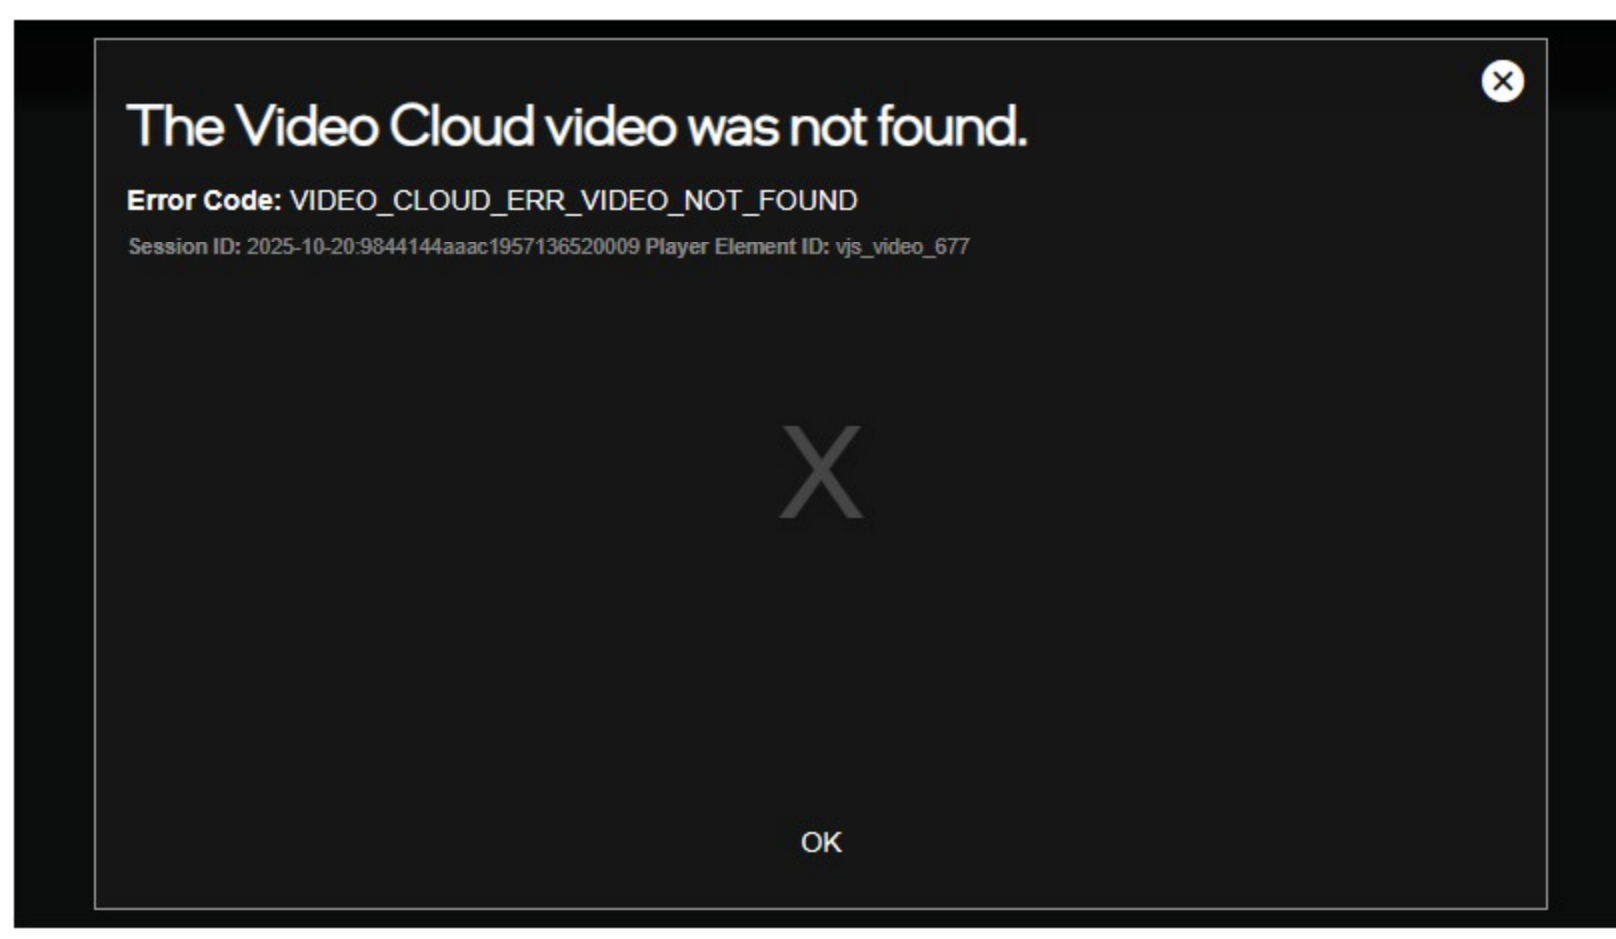
Neuromorphic visual scene understanding with resonator networks leverages the classic idea that brains solve perception through "analysis-by-synthesis." The research starts with a generative model, which is a probabilistic description of how a scene is constructed from components and their factors of variation, such as multiple objects in particular poses. Scene understanding must then infer the configuration of the generative model specific to the sensory input. In rich generative models, however, inference is computationally expensive, which has limited progress of this approach to date. The researchers propose inference with a resonator network, which can be efficiently executed on neuromorphic hardware. To achieve this, the input image is mapped to a high-dimensional but inherently structured vector representation. The information on individual factors of variation is held in different modules of the resonator network. A given scene can then be decomposed by the modular resonator network into vectors that describe different factors of variation of individual objects (such as shape, pose, and color).



A spiking neural network implementation of the resonator network mapped to Intel's Loihi neuromorphic research chip demonstrates the promise of neuromorphic hardware to accelerate such vector symbolic generative inference algorithms on resource-constrained edge devices. The Loihi implementation achieves 175x lower combined energy and latency per inference (also known as energy-delay product or EDP) compared to a conventional CPU implementation of the same resonator algorithm.

Visual Odometry in Robotics

Visual odometry with neuromorphic resonator networks extends the resonator network approach to address the practical robotics problem of using visual odometry to estimate a robot's motion and location from visual input. Unlike accelerometers, the tracking of self-motion through vision doesn't suffer from error accumulation. Researchers used an event camera attached to a robotics platform to transmit visual signals to the network (see cover image of June issue). A working memory of the visual environment is maintained by the network, and the resonator architecture is used to infer the pose of the camera.



The system outperforms deep learning approaches on standard VO benchmarks in both precision and efficiency — relying on less than 100,000 neurons without any training.

How Resonator Networks Work

While early in development, resonator networks are promising for enabling a myriad of applications. To interact in the real world, robotic agents need the ability to reliably understand and act within the surrounding environment. The generative modeling approach allows for understanding and provides transparency by linking sensory input to specific parameters of the model.

In a major innovation, the research shows how to extract factors of variations of objects using structured and decomposable vector representations. Resonator networks can explore the space of factor combinations in parallel using the strategy of "search in superposition." The network assesses many combinations of factors simultaneously in each iteration and will converge when it finds a combination of factors that sufficiently explain the input. Once parameters of the environment are successfully factorized or "disentangled," algorithms such as control, tracking, and navigation become straightforward.

Neuromorphic computing hardware, like Intel's Loihi, shows tremendous promise for AI applications using little energy, such as mobile, edge, and energy constrained devices. But to fully take advantage of these architectures, the algorithms must also be a good match. Intel's research efforts in the co-design of new types of neuromorphic algorithms and hardware, such as resonator networks, is at the core of a future where devices and machines will be capable of brain-like capabilities.

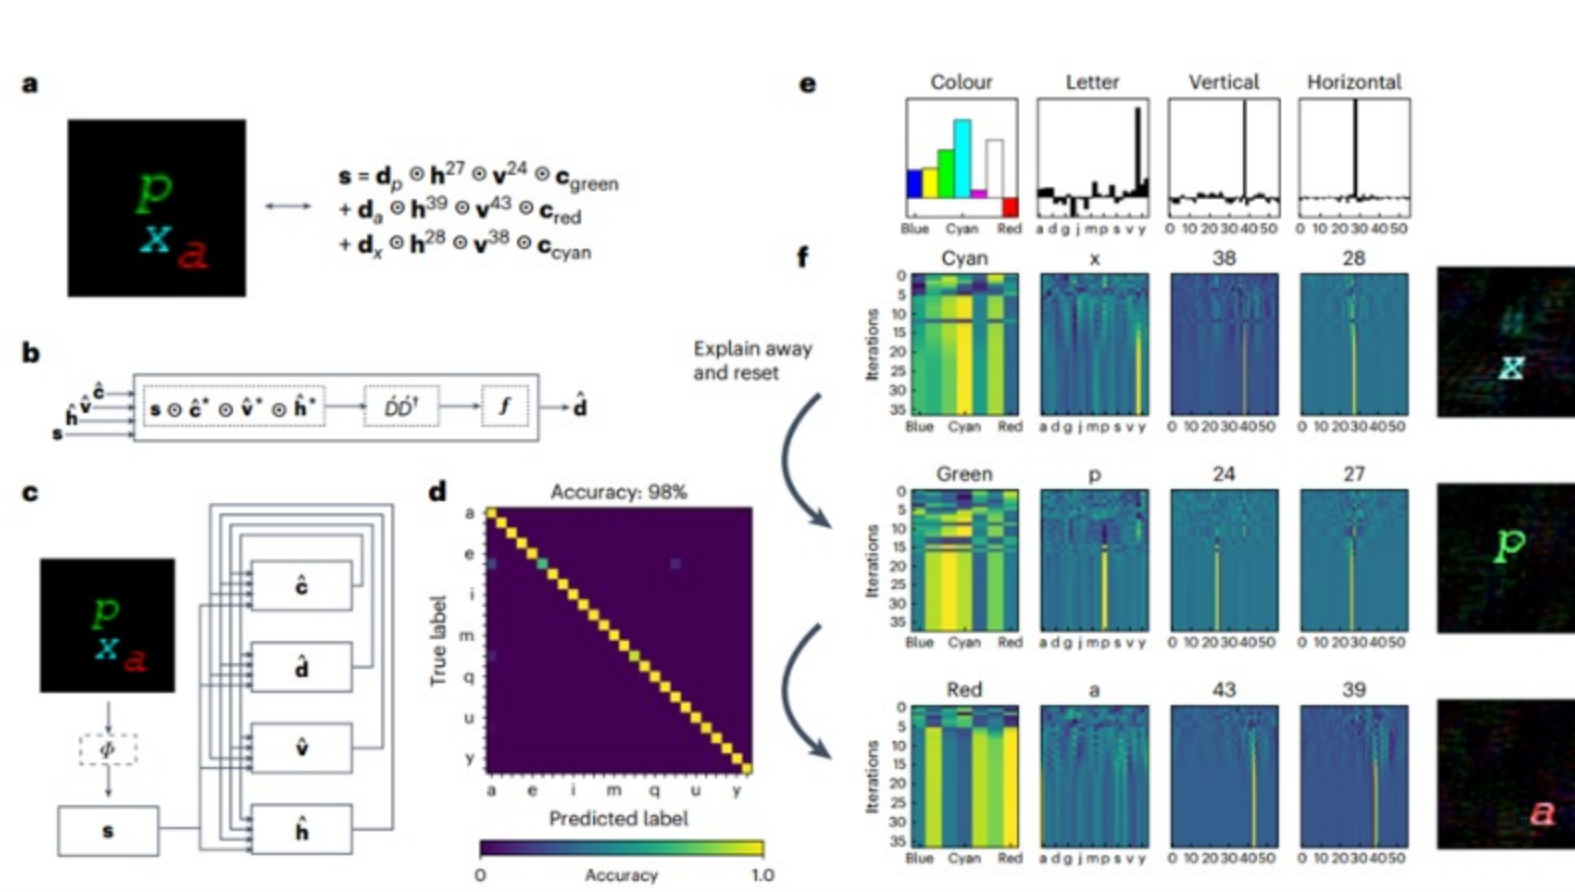


Figure 1. Resonator network for inferring shape, color, and translation. 1a. A synthetic scene and the structured vector representation, explaining the image objects in terms of the factors of variation — for a given pixel image, the task of the resonator is to recover this explanation. 1b. A resonator module consisting of binding stage, local memory, and neural nonlinearity. 1c. Image encoding and circuit diagram of the modules in the resonator network. 1d. Confusion matrix on translation benchmark task with a single object. The overall performance of the network is 98.4%. 1e. Confidence levels in each resonator module after converging to a factorization of object x. The maximum values are taken as the explanation of the object. 1f. Each row of panels shows the time evolution of the confidence values of values in the different factors, color, letter identity, x- and y-position. Confidence levels are depicted as a heatmap (same color scale as in d), with time evolving vertically down and different factor values represented horizontally. After several iterations, the resonator network converges to the explanation of an object (final explanation visualized in the far-right column) and remains stable. The final state in the first row corresponds to the bar graph in e). The explained object is removed from the input vector and after reset, the resonator network explains another object (rows 2 and 3). (Image credit: Nature Machine Intelligence).

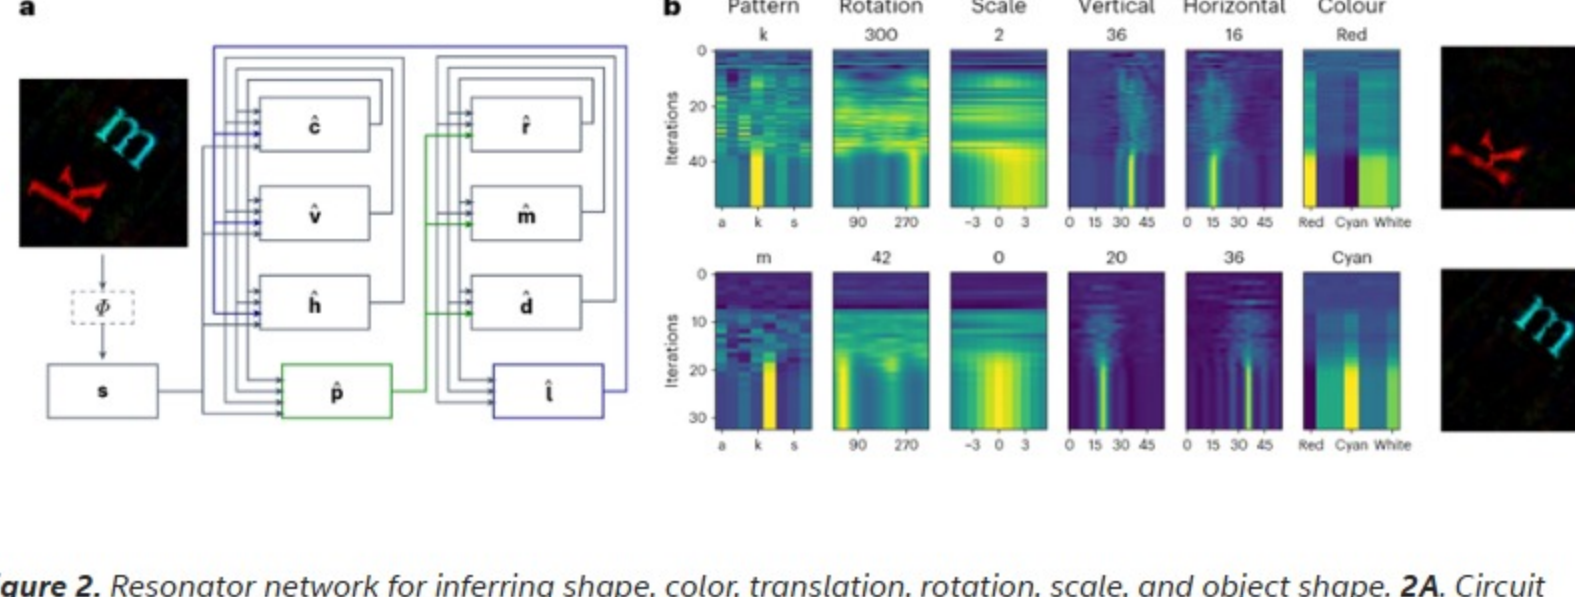


Figure 2. Resonator network for inferring shape, color, translation, rotation, scale, and object shape. 2A. Circuit diagram of hierarchical resonator, which alternates between modules in cartesian coordinates for the factors color, x- and y-position, and log-polar coordinates for the factors rotation angle, scale, and shape. 2B. Waterfall diagrams of evolution of confidence values during inference for explanation of letter k (top row), and letter m (bottom row). (Image credit: Nature Machine Intelligence).

Translate

Tags: AI Intel Labs neuromorphic

0 Kudos

About the Author



Scott Bair is a Senior Technical Creative Director for Intel Labs, chartered with growing awareness for Intel's leading-edge research activities, like AI, Neuromorphic Computing and Quantum Computing. Scott is responsible for driving marketing strategy, messaging, and asset creation for Intel Labs and its joint-research activities. In addition to his work at Intel, he has a passion for audio technology and is an active father of 5 children. Scott has over 23 years of experience in the computing industry bringing new products and technology to market. During his 15 years at Intel, he has worked in a variety of roles from R&D, architecture, strategic planning, product marketing, and technology evangelism. Scott has an undergraduate degree in Electrical and Computer Engineering and a Masters of Business Administration from Brigham Young University.

You must be a registered user to add a comment. If you've already registered, sign in. Otherwise, register and sign in.

Comment

Community support is provided Monday to Friday. Other contact methods are available here. Intel does not verify all solutions, including but not limited to any file transfers that may appear in this community. Accordingly, Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade. For more complete information about compiler optimizations, see our Optimization Notice.