The Squirrel Effect

Mathew M. Hubbard | October 29, 2026

Unexpected Results and the Value of Following the Data

While serving as a board member and consultant to an AI startup, I was recently reminded of the value of a strong, independent engineering team and the importance of following unexpected results wherever they lead.

The company was developing a custom-trained AI model designed to identify a specific target in an outdoor environment. The application required a high degree of identification confidence under a wide range of operating conditions, and considerable engineering effort had been invested in developing the model and preparing the training data.

Because the system would operate primarily outdoors, the training data also included a limited number of images representing other objects likely to appear within its field of view. Among these were squirrels.

The squirrels were something of an afterthought. They were not particularly important to the application, but were simply another object the system might encounter during normal operation. Consequently, comparatively little effort had been devoted to their inclusion in the training data.

The results of subsequent testing, however, revealed something unexpected.

When evaluated against random test samples that included the intended target, the model demonstrated an approximately 30 percent higher positive-identification rate for squirrels than for the target it had been specifically trained to recognize.

Considering the relative amount of training devoted to each object, the result was counterintuitive. The model was significantly more successful at identifying an incidental object than the one around which the application had been developed.

Understanding the Performance Difference

The results raised an obvious question. Why would a relatively small number of squirrel training images produce better identification performance than the much larger and more carefully developed training dataset for the intended target?

Improving model performance often involves expanding training datasets or refining the model itself. In this case, however, the more useful question was why an incidental object was already being identified so successfully.

The engineering team investigated the discrepancy between the two sets of training data.

The investigation identified a difference in the methodology used to crop the training images. Because the squirrel images had been added almost as an afterthought, the cropping rules used in their preparation differed from those established for the intended target. What initially appeared to be an inconsequential variation in data preparation turned out to have a significant effect on identification performance.

© 2026 Mathew M. Hubbard. All rights reserved. • mathewhubbard.com

Machine-learning performance depends on complex relationships within the training data, making the effects of individual data-preparation decisions difficult to predict in advance. Even a technically reasonable approach must ultimately be validated through testing.

The engineering team recognized that the squirrel-identification results challenged the assumptions underlying the original target-image preparation. Rather than dismissing the anomaly, they investigated its cause and applied what they learned to the intended target's training data. Additional testing verified the resulting performance improvement.

The specific technical details remain proprietary. However, the important observation is that a meaningful improvement in identification performance was achieved through changes to the training data rather than a fundamental redesign of the model itself.

The Value of Unexpected Success

What impressed me most when the team reported their findings was not simply the performance improvement, but how they had arrived at it. They recognized an unexpected result, questioned the underlying assumptions, and followed the evidence to a solution without needing to be directed to do so.

Engineering organizations routinely devote substantial resources to understanding failures. When a product or system does not meet its performance objectives, the response typically involves examining assumptions, isolating variables, and identifying the underlying cause.

Unexpected success does not always receive the same level of scrutiny.

When something performs better than anticipated, the result is generally welcomed. Unless the improvement is directly related to the primary development objective, however, there may be little incentive to investigate why it occurred.

This can represent a missed opportunity. An unexpected failure identifies something that does not work as intended, while an unexpected success may reveal an approach that works better than the one currently being used.

In this case, improving squirrel identification offered little direct benefit to the intended application. Nevertheless, the engineering team recognized that understanding the performance advantage might provide insights applicable to the broader system.

The Squirrel Effect

Throughout my career, I've worked with engineering and product teams developing technically sophisticated products for demanding applications. One of the characteristics I've come to value most is the willingness of engineers to challenge their own assumptions when the data suggests something unexpected. While development resources must naturally be concentrated on the primary objective, useful information can sometimes originate from areas receiving relatively little attention.

I think of this as the Squirrel Effect: the possibility that an unexpected result in a seemingly unimportant part of a system can reveal an improvement applicable to the system as a whole.

The underlying principle is straightforward. Unexpectedly good results deserve investigation just as much as unexpected failures.

© 2026 Mathew M. Hubbard. All rights reserved. • mathewhubbard.com

Sometimes you have to follow the data, no matter where it leads you.

Even if it leads you to a squirrel.

Online publication and PDF edition

Read the article online: https://mathewhubbard.com/articles/the-squirrel-effect.html

PDF edition: https://mathewhubbard.com/articles/pdf/the-squirrel-effect.pdf

© 2026 Mathew M. Hubbard. All rights reserved. • mathewhubbard.com