

Digital Twins: Definitions, Distinctions, and a Short Taxonomy

A. Author and B. Builder

2024 (synthetic sample)

Synthetic sample paper. This document was written solely as a sample source for the chitragupta documentation's worked examples. It imitates the shape of a survey paper; its claims are illustrative, not scholarship. Public domain (CC0).

Introduction

A digital twin is a virtual representation of a physical system that is kept consistent with that system for the length of its operating life. The phrase is often used loosely, and much of the early literature disagrees about where a simulation ends and a twin begins. This short overview fixes the vocabulary used across our sample corpus.

Three levels of integration

We distinguish three levels of integration between a physical object and its virtual counterpart. A **digital model** exchanges data with its physical counterpart only through manual steps: an engineer updates the model when the plant changes, and nothing flows back automatically. A **digital shadow** adds an automatic data flow in one direction, from the physical object to the virtual one, so the representation tracks reality without human effort. A **digital twin** closes the loop: data flows automatically in both directions, and a change in the virtual object can actuate a change in the physical one.

The distinction matters in practice because the engineering cost rises sharply at each level. In our synthetic survey of forty deployments, twenty-eight systems described as twins were, by the definition above, shadows: they monitored faithfully but actuated nothing.

What a twin is for

Four uses dominate. First, monitoring: the twin renders the current state of the asset legibly. Second, prediction: simulation over the twin's state estimates behaviour under

hypothetical loads before they occur. Third, optimisation: the closed actuation loop lets predicted improvements be applied automatically. Fourth, and least discussed, record-keeping: a twin's state history is an audit trail of the asset's life, valuable long after the asset is retired.

A note on fidelity

Fidelity is not free, and it is not always desirable. A twin that models every bolt of a wind turbine costs more to keep synchronised than the turbine costs to run. The useful question is not “how faithful is the twin” but “which decisions does the twin have to support”, and fidelity should be purchased decision by decision. In our sample deployments, the twins that survived beyond their pilot phase were uniformly the ones whose scope had been cut to a named decision list.

Conclusion

The model/shadow/twin distinction is a property of data flow, not of model sophistication. A crude model with two-way automatic flow is a twin; a exquisite model updated by hand is not. Keeping the vocabulary strict keeps expectations honest, and honest expectations are the difference between a deployed twin and an abandoned pilot.