Feature Dependence: A Method for Reconstructing Actual Causes in Engineering Failure Investigations*

Yafeng Wang

January 24, 2022

Abstract

Engineering failure investigations seek to reconstruct the actual causes of major engineering failures. The investigators need to establish the existence of certain past events and the actual causal relationships that these events bear to the failures in question. In this paper, I examine one method for reconstructing the actual causes of failure events, which I call "feature dependence". The basic idea of feature dependence is that some *features* of an event are informative about the features of its causes; therefore, the investigators can use the features of a known failure event to reconstruct details of its causes. I make explicit the structure of feature dependence and the evidential basis of its key premises, and show how feature dependence works in the investigation of the American Airlines Flight 191 accident.

Key Words: Actual Causation, Causal Reasoning, Historical Reconstruction, Features, Engineering Failure Investigations

1 Introduction

At 15:04 central daylight time, May 25, 1979, American Airlines Flight 191, a McDonnell-Douglas DC-10 airplane, crashed into a field less than a mile from the end of its departure runway at

*This is a preprint (the final draft before the peer review). Please cite the published Version: Wang, Y. (2022). Feature Dependence: A Method for Reconstructing Actual Causes in Engineering Failure Investigations. *Studies in History and Philosophy of Science*, 96:100-111. https://doi.org/10.1016/j.shpsa.2022.09.005

Chicago O'Hare International Airport. All 271 passengers and crew were killed, making this crash the second deadliest accident in the history of American commercial aviation. (NTSB 1979, 2)

Immediately following the accident, the National Transportation Safety Board (NTSB) conducted a crash investigation to uncover the events that led to the tragedy. After recovering the wreckage and interviewing witnesses, the investigators quickly discovered that a peculiar event had happened in the accident: As the airplane was about to lift off the runway, its left engine separated from the wing and landed on the runway. The investigators then raised a further question: What causal factors contributed to the separation of Flight 191's left engine? Answering this question led to discoveries of major maintenance flaws and procedural deficiencies that had plagued American Airlines and the entire DC-10 fleet in the U.S.

The NTSB's investigation of the crash of Flight 191 is an example of *engineering failure investigation*. The purpose of engineering failure investigations is to identify the causes of engineering system failures and to improve safety by adopting measures to prevent similar failures from occurring in the future. The problem of determining the causal factors that contributed to Flight 191's engine separation has the following characteristics:

1. The outcome is known but some causes of the outcome are not.
2. The outcome is a past event[1], and the problem is to determine its *actual* (token[2]) causes.
3. Solving the problem requires establishing both the existence of certain other past events and the causal relationships that these past events bear to the outcome.

Problems that have the above characteristics are problems of *reconstructing*[3] *the actual causes of past events*. Engineering failure investigators often encounter problems of this type in their investigations. The philosophical question I address in this paper is: What kinds of evidential reasoning do the engineering failure investigators use to reconstruct the actual causes of failure events?

There are a few discussions of the methodology of engineering failure investigations in the philosophy literature, but they are motivated by other philosophical questions and concerns. Galison

[1]Strictly speaking, the outcome could be a state of affair or a temporally extended and spatially dispersed sequence of events. Even though I focus mostly on spatially and temporally compact individual past events in this paper, my uses of the term "event" could be interpreted loosely to include both states of affairs and sequences of events.

[2]In this paper, I use the terms "actual causation" and "token causation" interchangeably.

[3]In this paper, the term "reconstructing" means "gaining epistemic access to".

2000 argues that there are "unsolvable tensions between competing norms of explanations" (Galison 2000, 4) in the NTSB's accident reports. One tension is between what he calls "localizing accounts" and "diffusing accounts" of accidents: The investigators could localize their causal explanations of the accident by tracing the causal chains backward to sites of technological failures and stop there. Or they could diffuse their causal explanations by continuing to trace the causal chains backward to include more human and institutional factors. Tensions like this make it difficult for the investigators to determine the scope and scale of their causal explanation of the accident.

Galison uses the NTSB's investigation of the United Flight 232 accident to illustrate his argument. On July 19, 1989, United Airlines Flight 232 suffered a catastrophic failure of its tail-mounted No.2 engine in-flight. The investigators found that the stage 1 fan disk of the No.2 engine disintegrated due to a microscopic fault in the disk bore. Tracing the causal chains backward from this fault led to flawed design philosophy, outdated manufacturing procedures, and shoddy maintenance practice. In principle, the investigators could further trace the causal chains backward to include factors about safety cultures, economic interests, and regulative policies (ibid., 38). At what point should the investigators stop expanding the scope of their causal explanations of the accident? Galison's essay highlights difficulties in answering this question.

More recently, Hanley 2021 proposes an account of causal selection in safety sciences. Hanley defines causal selection as the selection of particularly important causes among many causal factors of a phenomenon. Using the Bhopal Gas Tragedy as a case study, Hanley argues that pragmatic details help explain why safety scientists select systemic causes of sociotechnical failures as more important than proximate causes[4]: The primary pragmatic goals of safety scientists are predicting and preventing similar accidents in the future. Because of their higher causal stability and their influence on the stability of other causes, systemic causes are more efficient for predicting and preventing similar future accidents.

Both Galison 2000 and Hanley 2021 make valuable contributions to the methodology of engineering failure investigations, but neither work specifically addresses how investigators reconstruct the actual causes of failure events. Galison's essay raises important questions about which parts of the causal history of an accident should be included in the causal explanations of the accident,

[4]Examples of proximate causes include human errors and component failures, whereas examples systemic causes include factors relating to the overall design and organization of the sociotechnical systems.

but it lacks a philosophical account of how the investigators gain epistemic access to the causal history of the accident in the first place. Reconstructing actual causes is also different from causal selection. Causal selection involves picking out important causes from a set of *known* causes of a phenomenon. Therefore, causal selection presupposes that the investigators have already known a set of causes of the outcome, whereas causal *reconstruction* is precisely about getting to know what these causes are.

In this paper, I provide a partial[5] answer to the question of how engineering failure investigators reconstruct the actual causes of failure events. I identify *one* type of evidential reasoning commonly used in practice, which I call "feature dependence". The basic idea of feature dependence is that some of the *features* of an event is informative about the features of its causes; consequently, investigators can use the features of a known event to reconstruct the details of its causes. The primary goal of this paper is to make explicit the structure of feature dependence and the evidential basis of its key premises.

This paper is organized as follows. I begin by explaining the basic ideas of feature dependence using an example from the United Airlines Flight 585 accident investigation. Next, after making explicit a few assumptions about events, I introduce the concept of "features of an event" and illustrate three main types of features with examples. I then introduce a type of statement called *feature dependence statements*, which states that (some of) the features of an event depend on features of its causes. After discussing the basic forms and subtypes of feature dependence statements, I propose an account of the evidential basis of feature dependence statements, using the idea of *event-based causal systems*. Finally, I articulate the structure of a type of argument called "descriptive enrichment". Descriptive enrichment uses feature dependence statements as premises, and its basic idea is that we combine multiple features of an outcome to construct a detailed description of one of its causes. I show how descriptive enrichment works in practice via an example from the American Airlines Flight 191 investigation.

[5] The answer provided by this paper is partial because feature dependence is not the only type of evidential reasoning that can be used to reconstruct the actual causes of failure events.

2 Basic Ideas of Feature Dependence

The basic idea of feature dependence can be explained using an example from the investigation of the United Airlines Flight 585 crash. On March 3, 1991, United Airlines Flight 585, a Boeing 737-200 aircraft, crashed near Colorado Springs Municipal Airport in Colorado Springs, Colorado. After collecting witness reports and examining the radio data and data recorded in the two black boxes[6], the NTSB investigators arrived at the following picture of the last 14 seconds of the accident flight: United Flight 585 was in its final approach at about a thousand feet above the ground. When it was about to align with the runway, the aircraft momentarily leveled the wings and then rolled and yawed[7] to the right at an abnormally high rate. In the cockpit, the pilots reported emergency actions that would normally initiate a go-around on landing, but they did not succeed. The roll and yaw continued, and the airplane soon flipped over with its nose nearly straight down and dived towards the ground. (NTSB 2001, 3-4)

In aeronautics, a *loss of control* is defined as a significant and unintended departure of an aircraft from its intended flight path and flight envelope[8] (Jacobson 2010). The witness reports and the recorded data showed that Flight 585 suffered from a loss of control during critical moments of its final approach, but they provided no obvious answer as to *why* Flight 585 lost control.

The NTSB investigators examined many hypothetical causes of Flight 585's loss of control, one of which was an encounter with a mountain rotor. Rotors are turbulent vortices of air created by high winds blowing across mountainous terrain. The weather reports and witness reports of the day of the accident indicated the presence of rotors (NTSB 2001, 33). Consequently, the NTSB investigators examined the possibility that an encounter with a rotor caused the loss of control of the aircraft. Eventually, the investigators used the following argument against this possibility:

- Premise 1: **If** an encounter with a rotor caused Flight 585's loss of control, **then**, to account for the time history[9] of the aircraft's heading[10] during the loss of control, the rotor must have followed the flight path of the accident airplane for about 8 seconds and increased in strength

[6]The two black boxes are the flight data recorder (FDR), and the cockpit voice recorder (CVR).

[7]Yawing is the rotation of an aircraft about its vertical axis so that its nose moves from side to side.

[8]A flight envelope for a specific aircraft is defined as the limits of speed, altitude, and angle of attack that the aircraft cannot safely exceed.

[9]A time history is a sequence of values of a variable measured at a set of fixed times.

[10]The heading of an aircraft is the compass direction the aircraft's nose.

to about 1.8 radians (103 degrees) per second (NTSB 2001, 125-126).

- Premise 2: However, it was extremely unlikely for there to be a rotor that followed the flight path of United flight 585 for 8 seconds and increased in strength to about 1.8 radians per second (ibid., 131).

- Conclusion: It was extremely unlikely that United flight 585's loss of control was the result of an encounter with a rotor.

This argument is an instance of a *feature dependence argument*. Premise 1 is an instance of what I call *feature dependence statements*. To make explicit the structure of this statement, let event E be the loss of control of the aircraft during the final seconds of the flight. Let C be a hypothetical event, namely a (hypothetical) encounter with a mountain rotor. I represent Premise 1 schematically as follows:

- **If** the (hypothetical) event C was a cause of the (known) event E, **then**, given that event E had a feature F_E, event C must have a feature F_C.

I will say more about what I mean by "feature" in the next section. For now, features can be understood roughly as the details of an event. In the above schematic representation, the feature F_E consists of the time history of the aircraft's heading during the loss of control. This feature is a detail of the known event E (the loss of control of Flight 585), and the investigators had access to it via the data recorded by the flight data recorder (FDR). The feature F_C consists of the specific trajectory of the rotor (i.e., following the flight path for 8 seconds) and the specific strength of the rotor (increased to 1.8 radians per second) during the (hypothetical) encounter. This feature is supposed to be a detail of the hypothetical event C (the aircraft encountered a rotor). What Premise 1 accomplishes is an enrichment of the causal hypothesis that C was a cause of E: *Assuming* this causal hypothesis is true, Premise 1 enables us to infer more details about the hypothetical event C, namely that C had a feature F_C.

The main evidence for Premise 1 came from an analysis by Boeing engineers, who used aerodynamic modeling and computer simulations to examine the flight characteristics of Boeing 737-200 aircraft in a rotor-generated loss of control. Boeing developed flight control, aerodynamics, engine, and rotor models to derive force and moment time histories of a Boeing 737-200 aircraft, and equations of motion to convert these forces and moments into airplane motions. The inputs to the computer simulations included properties of the rotor (core radius, tangential velocity, position, and orientation); flight control inputs; engine thrust time history; and various initial conditions such as starting position, altitude, ground track angle, and ground speed. Outputs of the computer simulations included calculated ground track that could be compared with radar data; calculated airspeed, load factor, and heading time history that could be compared with FDR data, and calculated impact attitude that could be compared with crash site data (NTSB 2001, 56). Since the properties of the hypothetical rotor were unknown, the investigators adjusted the values of these variables to achieve matches between simulation results and recorded data on Flight 585. The simulations showed that to achieve a close match with the recorded heading data, a rotor strength of 1.8 radians per second would be required, and the aircraft would also have to remain in the rotor core for about 8 seconds (ibid., 130).[11]

Premise 2 is a rejection of the enriched causal hypothesis. In particular, Premise 2 rejects the statement that there was an event C with feature F_C. Part of the support for Premise 2 came from background knowledge: The strongest rotors ever documented in the Colorado Springs area had a strength of about 0.05 radians per second, which is much weaker than the hypothesized rotor of about 1.8 radians per second (ibid., 130). Another part of the support for Premise 2 came from the following argument: if United Flight 585 had penetrated the low-pressure core of a rotor of 1.8 radians per second and remained there for 8 seconds, its FDR would almost certainly have recorded signature changes in indicated airspeed and altitude (due to the low ambient pressure within the rotor), and its CVR would have recorded sounds characteristic of intense rotors. None of these traces could be found in the FDR and CVR data from the accident airplane (ibid., 131).

My purpose in describing the NTSB's argument in the United Flight 585 example is to display the structure of the feature dependence statement used in it. Later in this paper, I propose a general account of the evidential basis of this type of statement. Before doing that, I first explicate my uses of the terms "events" and "features".

[11] The accuracy of the simulations depended on the fidelity of the aerodynamic model of Flight 585's flight conditions during the accident. The model was validated by flight tests, but the validation process was limited by safety factors that restricted the flight tests being conducted (NTSB 2001, 86). Measurement errors and instrumentation calibration errors also affected the simulation accuracy. For instance, the directional gyroscope that provided heading information could have introduced errors when the aircraft was operating in certain flight attitudes, and the investigators had to develop a computer program specifically for correcting this type of error (ibid., 87).

3 Events and Features

Feature dependence is a type of evidential reasoning that helps determine the causes of a past *event*. Let me clarify my uses of the term "event". I make three major assumptions about events in this paper:

1. First, events occupy spatial and temporal regions.
2. Second, for each event, certain entities are "involved in" it.
3. Third, each event has many different *features*.

I will use the crash landing of United Airlines Flight 232 to illustrate the first two assumptions. When the tailed-mounted engine of United Airlines Flight 232 suffered a catastrophic failure in-flight, fragments from the failed engine damaged all the three hydraulic systems, causing the aircraft to lose all hydraulically powered flight controls. The pilots continued flying the aircraft by adjusting the power of the two remaining engines, and they eventually crash-landed the aircraft at Sioux Gateway Airport in Sioux City, Iowa, at about 16:00 central time (NTSB 1990, v).

First, as an event, the crash landing of United flight 232 occupies a particular spatial and temporal region. The event's temporal boundary began at the first ground contact of the aircraft by the right wing and ended when the final pieces of the aircraft came to rest. The spatial boundary of the event is more difficult to define precisely, but we know the event is located within Sioux Gateway Airport in Sioux City, Iowa. More specifically, the event occurred more or less along runway 22, where the aircraft crash-landed. The precise shape of the spatial boundary of the event is defined by how different parts of the aircraft broke apart and the parts' trajectories until they all came to rest.

Second, many entities were "involved in" the crash landing event. By "entities", I mean physical objects in general, including inanimate objects and living things. I do not have a precise account of the term "involved in", but it is easy to illustrate with examples. Some entities involved in the crash landing event already existed before the event. For instance, the right wingtip, the right main landing gear, the right engine, and all the passengers on board the flight were involved in the crash landing, and they all existed before this event. Some entities involved in the crash landing event ceased to exist after the event. For instance, the fuselage no longer existed after the crash, nor did the 111 people on board who died from the crash. Finally, some entities involved in the crash landing event were created during the event. For instance, the many fracture surfaces formed during the crash landing were new entities created during the event. All of the above entities were involved in the crash landing event.

The third assumption that events have features requires a more extended discussion. Intuitively, features of an event are details of that event. We can think of an event as a concrete being and each feature of the event as a particular aspect of this concrete being. Beyond that, I do not have a metaphysical account of what features are.[12] Fortunately, to understand how feature dependence works, a metaphysical account of features is unnecessary. Instead, it suffices to describe the main types of features commonly used in engineering failure investigations, which I will do in the rest of this section.

There are three main types of features commonly used in engineering failure investigations:

1. First, the properties of the entities involved in an event are features of that event.

2. Second, the identities of the entities involved in an event are features of that event.

3. Third, the spatial and temporal characteristics of an event are features of that event.

I do not claim that these three types of features exhaust all the possibilities of features, nor that the different types must be mutually exclusive. The only claim I make is that these three types of features are commonly used in causal reasoning about engineering failures.

Let us begin with the first type of features of an event. Suppose that the time duration of an event E is the interval $[t_1, t_2]$, and that an entity X was involved in the event E. Then the properties of X during the period $[t_1, t_2]$ are features of the event E. Similarly, if multiple entities $X_1, \ldots, X_n$ were involved in the event E, then the relations among $X_1, \ldots, X_n$ during $[t_1, t_2]$ are features of E.[13]

To illustrate this type of feature of an event, I discuss two examples. The first example concerns what I call "fracture events", and I argue that the properties of the fracture surfaces formed during

[12]My notion of "feature" has affinities with L.A. Paul's notion of "aspect" in "Aspect Causation" (Paul 2004). There are two main differences: First, Paul defines an aspect as a property instance, whereas my notion of feature is broader and arguably encompasses metaphysically heterogeneous categories. Second, my notion of feature is not motivated by metaphysical concerns, but rather by a methodological need to capture a pattern of evidential reasoning.

[13]For a property P of an entity X involved in an event E to be a feature of E, X must have property P at some point during the time span $[t_1, t_2]$ of E. For instance, if X only acquired property P after the event E was over, then P is not a feature of E.

a fracture event are features of that event. The second example concerns an aircraft encounter with a mountain rotor, and I argue that the properties of the aircraft, the properties of the rotor, and the relationships between the two during the counter are features of that event.

First, a fracture event is an event in which a fracture (crack) is formed.[14] As understood in failure analysis, fracture events typically consist of three phases: crack initiation, crack propagation, and final separation of the parts. The fracture surfaces were created during all three phases and were finally completed during the last phase. The fracture surfaces were new entities created during the fracture events, and they were involved in the fracture events. Hence the properties of the fracture surfaces created during the fracture events are features of the fracture events.

Some properties of fracture surfaces are macroscopic and can be observed using the naked eye. For instance, the size and shape of the fracture surfaces, the smoothness or ruggedness, and distortions and deformations on the crack surfaces are generally observable. In a particular type of fracture known as brittle fracture[15], the fracture surfaces often display a series of V-shaped marks known as the chevron marks. The chevron marks are informative because they point to both the fracture's origin and the direction of crack propagation.

Other properties of the fracture surfaces can only be observed using higher-resolution devices such as scanning electron microscopes. Consider a type of fracture known as fatigue fracture. Unlike overload fractures caused by the application of a single force greater than the strength of the materials, fatigue fractures generally occur after many stress cycles—think about how we break a paper clip by repeatedly bending it back and forth with each back-and-forth being one stress cycle. For certain types of materials (e.g., pure metals and some ductile alloys), when the fatigue fracture surfaces are observed under the scanning electron microscope, they often show a series of wave-like ridges called "striations". Striations are microscopic features of the fatigue fracture event, and research has indicated that each striation is the result of a single stress cycle. As a result, the number of striations is informative about the number of stress cycles that have caused fatigue fracture.

The above macroscopic and microscopic properties of fracture surfaces are all features of the

[14] Depending on the type of fractures in question, a fracture event can take a relatively short time or a very long time. Perhaps "crack growth events" will be a more appropriate term for describing them.

[15] Brittle fracture is the sudden, very rapid cracking of a part under stress where the material exhibited little or no evidence of deformation before the fracture occurs.

fracture event. However, not all current[16] properties of the fracture surfaces are features of the fracture event. Suppose an aircraft broke up in the air above an ocean, and the newly created fracture surfaces were immersed in the seawater for days before they were found. As a result, the fracture surfaces suffered from some corrosion damage. The corrosion damage is a current property of the fracture surfaces at the time of the investigation, but they are not features of the fracture events because they were created outside of the event's time span.

My second example is the encounter of an aircraft in flight with a mountain rotor, which, I repeat, is a type of turbulent vortices produced by high winds over a mountain range. By an encounter, I mean the aircraft and the rotor crossing paths, and the aircraft penetrating the low-pressure core of the rotor. The temporal duration of the encounter is determined by how long the aircraft remained in the rotor's low-pressure core. The aircraft and the rotor are entities involved in the encounter event.

Let us say that the time interval of the encounter is $[t_1, t_2]$. Since the rotor was involved in the event, the rotor's properties during this period are features of the encounter event. For instance, the rotational speeds of the rotor during $[t_1, t_2]$ is a feature of the encounter event. More precisely, since the rotational speeds of the rotor could constantly be changing during $[t_1, t_2]$, we could say that the rotational speed of the rotor at time t_i, for any t_i in the time interval $[t_1, t_2]$, is a feature of the encounter event. Similarly, the angle of interception (the angle at which the aircraft penetrated the low-pressure core of the rotor) is a relation between two entities involved in the encounter during the time span $[t_1, t_2]$. Thus it is also a feature of the encounter event.

This concludes my illustration of the first type of features of an event, namely the properties of the entities involved in the event.

The second type of feature of an event is the identities of the entities involved in the event. My conception of "identity" here is metaphysically minimal: To know the identities of entities involved in an event is to know *which* entities are involved in that event.

The importance of the identities of those entities involved in engineering failures is clear. Investigators are often interested in knowing *which* entities are involved in the *causes* of a given failure, partly for epistemic reasons: Identifying the entities involved in the causes of a failure event often makes it easier to infer additional information about how the failure event occurred. In general,

[16] The term "current" means "existing at the time of the investigation".

knowing one feature of an event can be informative about other features of the same event, and the identities of the involved entities can be a particularly informative type of feature.

It is straightforward to find engineering examples illustrating identities as features of an event. Recall that United Airlines 232 suffered from a catastrophic failure of its tail-mounted No.2 engine on July 19, 1989. Earlier in the investigation, the investigators suspected that some rotating part of the No.2 engine flew apart, and the fragments severed the hydraulic lines in the right horizontal stabilizer. What they wanted to know is *which* rotating parts of the No.2 engine severed the hydraulic lines. In this example, the identities of the No.2 engine parts that severed the lines—which turned out to be fragments from the stage 1 fan disk—are features of the hydraulic line severing event.

The third main type of feature of an event is the spatial and temporal characteristics of the event. First, the starting time point and endpoint of the event, and the time interval between them, are features of the event. For instance, in the case of a fatigue crack (created over a long time), the investigators often want to know when the crack was initiated. That is the starting time of the fatigue cracking event and is a feature of the event. In the example of an aircraft encounter with a rotor, the duration of the event (the amount of time the aircraft is in the rotor's low-pressure core) is a feature of it.

Second, for a complex event E composed of smaller "subevents", the time history of the subevents are features of the event E. For instance, the time history of the formation of microscopic striations is a feature of the fatigue crack growth event. Moreover, if a property of an involved entity constantly changes during the event, then the time history of the instances of that property are features of the event. For instance, the growth rate of a crack is a feature of the crack growth event.

Finally, the spatial characteristics of the event—including but not limited to the location of the event—are also features of the event. For instance, the location of the fracture on a piece of equipment, the location of the fracture origin, and the direction of the crack growth are all spatial features of the fracture event.

In sum, the properties of the entities involved in the event, the identities of the entities involved in the event, and the spatial and temporal characteristics of the event are three main types of features of the event. There does not appear to be an existing metaphysical category that covers all three types of features, which is why I use a metaphysically neutral term "feature" to characterize them all.

4 Feature Dependence Statements

Feature dependence is, first and foremost, a type of *argument*. Feature dependence arguments make use of a special type of premises, which I call feature dependence *statements*. In this section, I introduce the basic forms of feature dependence statements.

The basic idea of a feature dependence statement can be captured by the following conditional statement: "If an event C is a cause of another event E, then certain features of E depend on certain features of C." The antecedent of the conditional is a causal statement, which describes a causal relationship between two events. Moreover, in the next section, I will further argue that the consequent of the conditional is *also* a causal statement. More specifically, the consequent describes a causal relationship between two features, and the features can be represented as the values of two variables. In short, a feature dependence statement is a conditional statement that encapsulates an inference from one causal statement to another causal statement.

Next, I make a distinction between two types of feature dependence statements, namely *feature necessity* and *feature sufficiency*. The basic idea of feature necessity can be characterized as: "If an event C is a cause of an event E, then, given that E has certain features, C must have certain corresponding features." That is, certain features of the outcome *require* certain features of one of its causes. In contrast, the basic idea of feature sufficiency can be characterized as: "If an event C is a cause of an event E, then, given that C has certain features, E must have certain corresponding features." That is, certain features of a cause *imply* certain features of the outcome.

When the investigators reconstruct the causes of past events, they typically know more about the outcome than about its causes. This makes feature necessity a natural fit for reconstructing the causes of past events. Feature sufficiency also has its uses; e.g., it could be used in a hypothetical-deductive manner to test specific hypotheses about the features of past events. In this paper, however, I focus exclusively on feature necessity, mostly because of its central role in a type of feature dependence argument that I call "descriptive enrichment". Below, I provide two formulations of feature necessity, both of which appear in later discussions of descriptive enrichment.

- Formulation 1 of feature necessity: If an event E has a feature F_E, then, there must exist a cause C of E, such that C has a feature F_C.[17]

This formulation can be used when we know nothing about the hypothetical event C yet, and we need to infer its *type* first. Consider an example: "If the fracture surface created in the fracture event has chevron marks on it, then the fracture must have been caused by an overload."[18] In this example, event E is the fracture event; feature F_E is a property of fracture surfaces created in the fracture event, namely the chevron marks on the fracture surfaces; and feature F_C is the *type* of the cause C, namely an overload.[19]

- Formulation 2 of feature necessity: If an event C of type X is a cause of an event E of type Y, then, given that E has a feature F_E, C must have a feature F_C.[20]

This formulation is often used when we already know (or assume) some partial information about the (hypothetical) cause C of event E, and we want to infer *more* information about C. This formulation of feature necessity is useful for reconstructing actual causes because having more information about the putative cause C amounts to enriching a causal hypothesis. An enriched causal hypothesis can sometimes be easier to test, as we saw in the United Flight 585 example.

Formulation 2 of feature necessity is a type of conditional statement with a complex structure. How would investigators support this type of statement with evidence? I address this question in the next section.

5 Evidential Basis of Feature Dependence Statements

The second formulation of feature necessity states: "If an event C of type X is a cause of an event E of type Y, then, given that E has a feature F_E, C must have a feature F_C". This conditional statement encapsulates an inference from a causal relationship between event C and event E, to a dependence relationship between feature F_C of C and feature F_E of E. In this section, I address two main questions: First, what is the nature of the dependence relationship between feature F_C

[17] I choose to formulate the feature dependence statements in the present tense.

[18] An overload is an application of stress that exceeded the strength of the material.

[19] The type of an event is also a feature of that event. Arguably, the type of an event is reducible to (or supervenient on) other features of that event, including the three types of features I discussed earlier.

[20] The logical form of this statement is: $\forall C \forall E((\text{TypeX}(C)\ \&\ \text{TypeY}(E)\ \&\ \text{Cause}(C, E)) \rightarrow (F_E(E) \rightarrow F_C(C)))$.

and feature F_E? Second, what is the evidential basis of such a conditional? For the first question, I propose that the dependence relationship between the two features is *also* a causal relationship. For the second question, I propose that the inference can be supported by suitable background knowledge concerning certain *event-based causal systems*.

The antecedent of the second formulation of feature necessity specifies an actual or token-level causal relationship between two events. There are many philosophical analyses of actual causal relationships among events[21], motivated primarily by a set of puzzles including cases of redundant causation, causation by omission, and apparent failures of transitivity[22]. In this paper, I am not committed to a particular philosophical analysis of actual causal relationships among events, because my account of the evidential basis of feature dependence does not require it.

The consequent of the second formulation of feature necessity specifies how a feature of C makes a difference to a feature of E. To see this more clearly, we can rewrite the second formulation of feature necessity as follows: "If an event C of type X is a cause of an event E of type Y, then, if event C had not had a feature F_C, event E would not have had a feature F_E." The consequent of this rewritten version is a counterfactual conditional, which makes explicit how the feature F_C makes a difference to the feature of F_E. I claim that this difference-making relationship is also causal: It is a causal relationship between two features.

The evidential basis of the second formulation of feature necessity consists of background knowledge that justifies the transition from event causation to feature causation. What kind of background knowledge can justify this transition? I propose that the background knowledge about an *event-based causal system* can play such a justification role.

The easiest way to understand the idea of an event-based causal system is through an example. Consider a fractured component. Its very existence tells us that a "fracture event" E had occurred. Suppose that E is a fatigue fracture. To say that it is a fatigue fracture is to say that E was caused by a cause C of a certain type—the application of cyclic loading at a stress level less than the strength of the material.

Now, *if* we know that the fracture was caused by this kind of cyclic loading, then our background knowledge in material engineering tells us that a certain causal system S—call it fatigue—was at

[21]See Lewis 1974, Lewis 2004, Hall 2004, Halpern and Pearl 2005, among others.
[22]See Paul and Hall 2013 for a survey of some of these puzzles.

work. I call the fatigue causal system *event-based*, because it is instantiated *whenever* an event of a certain type (cyclic loading at a stress level lower than the material's strength) caused another event of a certain type (crack growth).

In general, an event-based causal system S has the following characteristics:

1. S is instantiated whenever an event of type X causes an event of type Y.

2. S supports multiple systematic dependence relationships between features of event type X and features of event type Y. To be more precise, each systematic dependence relationship here is a *type-level* relationship between types of features.

3. Each systematic dependence relationship between types of features is a *causal* relationship. Moreover, the causal relationships among features can be analyzed using James Woodward's interventionist framework.[23] For instance, the relata of the causal relationships are variables, with each variable representing a type of feature. Moreover, each causal relationship characterizes how features of event type Y would change in response to interventions on features of event type X when some other parameters of the causal system S are held fixed.

To illustrate how an event-based causal system can support multiple causal relationships connecting features of one event (type) and features of another event (type), consider the fatigue causal system again. I mentioned earlier that the fatigue causal system is event-based because it is instantiated whenever one type of event (cyclic loading of a certain kind) causes another type of event (crack growth). The study of the fatigue causal system is a part of an engineering field called *fracture mechanics*. Given our current background knowledge about fracture mechanics, knowing that a fatigue causal system is instantiated gives us a lot of information about how the features of the cyclic loading make a difference to the features of crack growth. Below are two examples of type-level causal relationships between features in the fatigue causal system.

The first example is a type-level causal relationship between the number of stress cycles and the number of striations on the fracture surfaces. Some materials such as aluminum alloys display

[23]See Woodward 2003. A consequence of analyzing causal relationships among features using the interventionist framework is that the second type of features I discussed in section 3—the identities of entities involved in an event—does not enter into feature dependence relationships. This does not mean that the identities play no evidential role in reconstructing failures. Sometimes, knowing the identity of an entity involved in an event allows us to infer other features of the *same* event. We will see an example illustrating this point in the section on descriptive enrichment: After the investigators determined that the identity of the object that struck the upper surface of the flange was the wing clevis, they were able to significantly narrow down the possibilities for the *timing* of the strike.

well-defined fatigue striations on the fracture surfaces, with each striation corresponding to one advancement of the fatigue crack front. Moreover, research showed that given appropriate background conditions, there is a nearly one-to-one correspondence between the number of stress cycles and the number of striations on the fracture surfaces. Here, the number of stress cycles is a feature of the cyclic loading event; the number of striations on the fracture surface is a feature of the fracture event; there is a systematic relationship between the two types of features. The systematic relationship is a type-level causal relationship because experimental interventions on the number of stress cycles would make a difference to the number of fatigue striations.

The second example is a type-level causal relationship between the stress intensity factor range and the crack propagation rate. Before introducing this example, I will first explain a fracture mechanics concept known as the stress intensity factor K. When we apply load to a part, sometimes the stress is not concentrated uniformly all over the part. When there is a crack already, stress is concentrated more on the crack tip. The stress intensity factor K represents stress intensity near the tip of a crack. The magnitude of K depends on the geometry of the part, size and location of the crack, magnitude of the load on the entire part, and how the load is distributed. For instance, K is proportional to the square root of the crack length. Assuming that the applied loading remained constant, the longer the existing crack is, the greater the stress intensity is at the crack tip.

In the fatigue causal system, a fluctuating stress intensity drives the crack to grow at some rate. Let the *stress intensity range* be the difference between the maximum and minimum stresses in a stress cycle. When a stress intensity range ΔK is applied to a material for some number of cycles ΔN, this drives the crack to grow in length by a specific amount Δa. Here ΔK and ΔN are features of the cyclic loading event, Δa is a feature of the crack growth event, and there is a systematic dependence of Δa on ΔK and ΔN. Again, the systematic dependence relationship is causal.

A more standard way of understanding the systematic relationship between Δa, ΔK, and ΔN is to specify how the growth rate of the crack depends on the stress intensity factor range. The growth rate of the crack length relative to the number of crack cycles is given by the ratio $\Delta a/\Delta N$. Using a continuous expression, the crack growth rate is the derivative da/dN. For a given material at a stress ratio[24], the crack growth rate da/dN and the stress intensity range ΔK can be plotted

[24]The stress ratio R is the ratio between the minimum and the maximum stress amplitudes in the cycle.

on a log-log scale. On the log-log plot, there is typically a straight-line region of da/dN over ΔK values. This region corresponds to the crack-propagation phase and is defined by the Paris Equation $da/dN = A(\Delta K)^m$, where A and m are constants depending on the material in question.[25]

In short, the Paris Equation defines a systematic relationship between the stress intensity factor range ΔK at a given time, and the crack growth rate da/dN at that time. If we think of the stress intensity factor range at a given time[26] as a feature of the cyclic loading event, and the crack growth rate at that time as a feature of the crack growth event, then the Paris Equation defines a type-level causal relationship between a feature of the cyclic loading event and a feature of the crack growth event.

The examples of the striation count and the Paris Equation show that the very same event-based causal system—fatigue—can support multiple type-level causal relationships between features of one event (cyclic loading) and features of another event (crack growth). All these type-level causal relationships among features are different aspects of the *same* causal system.

The above discussions about event-based causal systems also suggest a way of justifying feature necessity statements. Let us return to the second formulation of feature necessity: "If an event C of type X is a cause of an event E of type Y, then, given that E has a feature F_E, C must have a feature F_C." The following line of reasoning justifies the transition from the antecedent to the consequent of this conditional:

1. Suppose an event C of type X is a cause of an event E of type Y.
2. We have background knowledge about an event-based causal system S, which instantiates whenever an event of type X causes an event of type Y.
3. The event-based causal system S supports a type-level causal relationship R. Knowledge about R is part of our background knowledge about the event-based causal system S.
4. The dependence relationship between feature F_C and feature F_E is an instance of a type-level causal relationship R.

[25]To determine the values of these constants, engineers test multiple samples of each alloy at different temperatures, stress ratios, and atmospheric environments. They measure crack length versus the number of cycles for different levels of applied loading in the laboratory. The way Paris Equations are established fits well with the interventionist analysis of causation.

[26]More precisely, ΔK is the stress intensity factor range in a given stress cycle. Similarly for da/dN.

5. Finally, to justify feature *necessity*, the type level relationship R must have an additional characteristic: When the dependent variable takes on the value F_E, the independent variable can only take on the value F_C. To use interventionist terminology, the cause C has some degree of *specificity*.[27]

The right background knowledge is required for this chain of reasoning to go through: We do not always have sufficient knowledge about the event-based causal system that instantiates when an event of one type causes an event of another type. However, if we do have sufficient knowledge about the underlying causal system, we can justify a feature necessity statement by citing one known aspect of the system, namely one known causal relationship within the system that connects features of the relevant events.

There is, however, a major missing piece in my account of the evidential basis of feature necessity statements so far. I have argued that the justification of a feature necessity statement relies on our background knowledge about the presence of some event-based causal system S and some type-level causal relationship R supported by S. What I have neglected to mention is that the justification of the feature necessity statement also requires an *assumption*, namely that no other causal system S^* that can interact with S and modify the causal relationship R is present.

To see why such an assumption is necessary, consider the fatigue example again. Let event C be a cyclic loading (of an appropriate sort), and event E be a crack growth. We know that whenever cyclic loading of the appropriate sort causes crack growth, the fatigue causal system S is instantiated. According to some type-level causal relationship R supported by S, if event E has feature F_E, then event C must have feature F_C. For instance, F_C could be the number of stress cycles and F_E be the number of striations on the crack surface, or perhaps F_C could be the stress intensity factor range at a given time and F_E be the crack growth rate at that time. So far, these considerations support the feature necessity statement: "If an event C of type X is a cause of an event E of type Y, then, given that E has a feature F_E, C must have a feature F_C."

Now, however, suppose further that another event D—some type of corrosion—is also a cause of event E. We know that corrosion can obliterate striations, which changes the fidelity of striations

[27] For more on the concept of specificity, see Woodward 2010. Woodward takes causal relata to be variables capable of taking on many values. A cause is more specific concerning an effect if the mapping between the values of the cause and effect variables is closer to a bijection.

count. So, if both fatigue and corrosion were causes of crack growth, the one-to-one correspondence between striation count and cycles of loading would no longer be reliable. In such a case, a feature necessity statement that infers the cycles of loading from the striation count would not be justified.

Moreover, the combination of fatigue and corrosion also modifies the Paris Equation, which defines a relationship between the stress intensity factor range and the crack growth rate. The fatigue crack growth rate is enhanced by corrosion given the same level of stress intensity range, and this effect is seen in all three phases of crack growth. The specific effects of corrosion on fatigue crack growth depend on a variety of factors, including the type of environment, fatigue load levels, and corrosion type. (Brooks et al. 2003)

In short, the effect of corrosion on crack growth can be regarded as an event-based causal system S^* in itself, and S^* interacts with the fatigue causal system S and modifies the quantitative form of the causal relationship R supported by S. If we have reason to believe that corrosion is a causal factor of the subcritical crack growth, then we are no longer justified in using the original feature necessity statement to infer features of the cyclic loading event. Instead, we should use the following, more complicated feature necessity statement: (Here X is cyclic loading, Z is corrosion, and Y is crack growth.)

> If an event C of type X is a cause of an event E of type Y, and if an event D of type Z is another cause of event E, then, given that event E has feature F_E, event C must have feature F_C^*.

The evidential basis of this more complicated feature necessity statement is our background knowledge about how *two* event-based causal systems—namely, fatigue and corrosion in this case—interact with each other to influence crack growth. Of course, the statement still assumes that no *further* causal system that can modify the dependence relationship between feature F_C^* and feature F_E is at work.

In sum, part of the evidential basis of a feature necessity statement is evidence for the assumption that no other event-based causal system that can change the feature necessity relationship we are interested in is at work. If there is evidence *against* this assumption, then we need to modify the feature necessity statement, by considering how the multiple event-based causal systems at work interact with each other to influence features of a given outcome.

6 Descriptive Enrichment

Previously, my discussions have focused on the general forms and the evidential basis of feature dependence statements. In the rest of this paper, I will discuss how feature dependence statements can be *used* in a type of evidential reasoning that I call "descriptive enrichment".

Suppose we know that an event E had occurred with certain features, but we do not know what had caused E. The basic idea of descriptive enrichment is that we can use *multiple* features of the known event E to construct a description of *one* of the causes of that event. It works as follows:

- Premise 1: Event E is of type Y and has features F^0_E, F^1_E, F^2_E, etc.
- Premise 2: If event E has a feature F^0_E, then, there must exist a cause of E that is of type X.
- Conclusion 1: There exists a cause of E that is of type X.
- Premise 3: If an event of type X is a cause of E (which is of type Y), then, given that E has a feature F^1_E, this event of type X must have a feature F^1_X.
- Conclusion 2: There exists a cause of E that is of type X and has a feature F^1_X.
- Premise 4: If an event of type X is a cause of E (which is of type Y), then, given that E has a feature F^2_E, this event of type X must have a feature F^2_X.
- Conclusion 3: There exists a cause of E that is of type X and has features F^1_X and F^2_X.
- Continue the process to learn more features about this cause of E that is of type X, until we can crosscheck that such an event existed based on the features it is presumed to have.

Premise 2 is an instance of the first formulation of feature necessity. This formulation is used when we know nothing about the cause yet and we need to infer its *type* first. Premise 3 is an instance of the second formulation of feature necessity.[28] This formulation can be used when we already have some partial information about the (hypothetical) event of type X, and we want to infer more information about it. My justification of Premise 3 should be familiar by this point: The type Y of the event E, together with the type X of one of its causes, helps to determine an event-based causal

[28] In my previous formulations of feature necessity, I used the schematic letter C to denote the hypothetical cause of event E. I omitted this schematic letter in Premise 2 and 3 to highlight the following: Our epistemic access to this hypothetical event is mediated by a *description* of the form "there exists a cause of E with feature F_1, F_2, etc".

system S that connects the two events. Our knowledge about S tells us *which* features $F_E^i, F_E^j, \ldots$ of E depend on features of this cause, and we can choose to use these features $F_E^i, F_E^j, \ldots$ in the descriptive enrichment argument to infer more features of this cause.

To illustrate how descriptive enrichment works in practice, I describe an example of it in the investigation of the American Airlines Flight 191 accident. In this accident, the airplane's left (No.1) engine separated from the wing as the plane was lifting off the runway. By careful examination of the wreckage, the NTSB investigators established that the No.1 engine's separation originated with a fracture of the upper flange of the pylon aft bulkhead. (NTSB 1979, 48)

A pylon is a structure designed to carry the engine, and it is connected to the wing through a few spherical bearings. Two of the spherical bearings are aligned vertically in a pylon forward bulkhead—an upright front wall attached to the forward portion to the wing. A third spherical bearing is in the pylon aft bulkhead—a smaller, upright back wall attached to a clevis fitting on the wing's underside. The pylon aft bulkhead has a series of projecting rims around its periphery; these projecting rims are called "flanges", and they help to connect the bulkhead to the pylon itself. In the Flight 191 accident, the upper side of the flange of the pylon aft bulkhead fractured, which caused a series of further fractures that ultimately resulted in the separation of the engine from the wing. (ibid., 12)

The largest part of the fracture on the upper flange of the pylon aft bulkhead was about 10 inches long. Let the event E be the fracture event in which this 10 inches long fracture came into being. To identify what caused E, the investigators used a descriptive enrichment argument.

First, the investigators carefully examined the fracture surfaces and found that the marks on the fracture surfaces (e.g., the presence of chevron and tear marks) were all typical of overload. (ibid., 12) We can represent this discovery as the first step of the descriptive enrichment:

- Premise 1: If event E had a feature F_E^0 (the presence of chevron and tear marks on fracture surfaces), then, there must exist a cause of E that is of type X (is an overload[29]).

Second, the investigators identified a few more properties of the fracture surfaces that were informative: The *direction* of the chevron and tear marks on the fracture indicated that the rupture progressed downward at the center of the flange. The bottom portion of the fracture exhibited

[29] That is, it is an application of stress greater than the strength of the upper flange.

smearing consistent with the compression portion of a bending fracture. The smear was more prevalent in the center section of the upper flange structure but became less prevalent at the fracture's outer ends. (NTSB 1979, 12) Based on these features F_E^1 of the fracture event E, the investigators were able to infer another feature of the overload event, namely that it is downward bending stress applied at the upper surface of the center section of the upper flange. In short:

- Premise 2: If an event of type X (is an overload of the upper flange) was a cause of E (the fracture event that created the 10-inch fracture on the upper flange), then, given that E had a feature F_E^1 (more precisely, a few features including the spatial characteristics of the chevron marks, the smear on the bottom portion of the fracture, etc), this event of type X must have a feature F_X^1 (downward bending stress applied at the center section of the upper flange).

- Conclusion 1: Therefore, there existed a cause of E that is of type X, and it has a feature F_X^1. That is, the 10-inch long fracture was caused by the application of downward bending stress at the center section of the flange, which exceeded the flange's strength.

Merely knowing that the cause of the fracture is a downward bending overload, however, was still not enough. Many things could have imposed such stress, and the event could have happened at any point during the airplane's service history. To further narrow the space of possibilities, the investigators zoomed in on another feature of the fracture event. At the center of the fracture surface, there was a crescent-shaped deformation. Moreover, the shape of the deformation exactly matched the radius of the bottom surface of the wing clevis to which the aft bulkhead was paired. This was strong evidence that the clevis was the object that imposed the downward bending stress on the center section of the upper flange. (ibid., 12-18)

To represent this step of the argument more formally, I will break it down into two smaller substeps. Let the feature F_E^2 be the specific shape of the crescent-shaped deformation, the feature F_X^2 be the specific shape of the object that applied the downward bending stress to the upper flange, and the feature F_X^3 be the identity of the object that applied the downward bending stress to the upper flange, namely the wing clevis. Here F_E^2 is a feature of the fracture event E, while F_X^2 and F_X^3 are features of the overload event. The first substep can be represented by the following premise:

- Premise 3: If an event of type X was a cause of E, then, given that E had another feature F_E^2, this event of type X must have a feature F_X^2.

In plain English, Premise 3 says: "If the application of downward bending stress on the upper flange was a cause of the 10-inch fracture, then, given that the fracture surface contains a crescent-shaped deformation of a specific shape, the object that applied the downward bending stress to the upper flange must have a matching shape."

The second substep is an inference from F_X^2 to F_X^3. The inference is through elimination: Among all the reasonable alternatives, the only object that has the precise shape required by feature F_X^2 is the bottom of the wing clevis, so the clevis must be the object that applied the downward bending stress (i.e., the overload must have feature F_X^3). This second substep can be represented by the next premise:

- Premise 4: If this event of type X had feature F_X^2, it must also have feature F_X^3.

- Conclusion 2: Therefore, there existed a cause of E of type X, with features F_X^1, F_X^2, F_X^3. That is, the 10-inch fracture on the upper flange was caused by the bottom of the clevis applying downward bending stress to the upper flange and overstressing it.

Next, the investigators moved on to determine the *time* when the wing clevis had impacted the pylon aft bulkhead's upper flange. The investigators noticed that with the bulkhead to clevis attaching hardware in place, the upper surface of the flange was about 0.5 inches below the clevis's bottom. (NTSB 1979, 18) Therefore, for the clevis to have contacted the flange and deformed it, the attaching hardware through the clevis and the bulkhead's spherical bearing must have been removed. Moreover, the investigators further established that the attaching hardware remained in place throughout the accident. The attaching hardware could only have been removed when the pylon was installed or removed from the wing (i.e., during a maintenance operation). Consequently, the clevis must have contacted the upper flange during a maintenance operation. (ibid., 49)

To represent this step of the reasoning, let feature F_X^4 be the timing of the overload event, namely that it had occurred during a maintenance operation in which the clevis to bulkhead attaching hardware was removed. What this step says is that, given that the overload event consists of the clevis applying downward bending stress on the upper flange, and given that the clevis could only come into contact with the flange during a maintenance operation, it follows that the overload event occurred in a maintenance operation. More formally:

- Premise 5: If an event of type X was a cause of the fracture event E, and if this event had F_X^1, F_X^2, F_X^3, then this event must also have feature F_X^4.
- Conclusion 3: Therefore, there existed a cause of event E of type X, and it had features $F_X^1, F_X^2, F_X^3, F_X^4$. That is, the fracture event was caused by the bottom of the clevis applying downward bending stress to the upper flange during a maintenance operation in which the clevis to bulkhead attaching hardware was removed.

By this point, the investigators had completed the descriptive enrichment argument. They had arrived at a sufficiently detailed description of one cause of the fracture event, namely "the bottom of the clevis applying downward bending stress to the upper flange during a maintenance operation in which the clevis to bulkhead attaching hardware was removed". The description was sufficiently detailed because the investigators could check it against the maintenance records to see whether any actual event on record fits the description.

After examining the airplane's maintenance history documentation, the investigators found evidence that there was an actual event that fits the description they constructed. On March 29-30, about eight weeks before the crash, the accident aircraft went through a pylon removal procedure at an American Airlines facility in Tulsa, Oklahoma. The record showed that the midnight shift started the operation by removing the hardware that attached the aft bulkhead to the clevis before going off duty. Since the forward bulkhead attachment hardware was still in place, it could act as a pivot in the case of any inadvertent loss of forklift support to the engine and pylon assembly, resulting in an upward movement at the aft bulkhead's upper flange and bringing it into contact with the wing clevis. (NTSB 1979, 49)

The next day, two mechanics who performed the maintenance procedure stated that they saw the upper lug of the aft bulkhead resting against the bolts attaching the wing clevis to the wing. What they observed and reported was a spatial configuration consisting of the aft bulkhead resting against the clevis in a specific way. Post-accident tests showed that such a spatial configuration could only have occurred if the clevis had deformed the upper flange by a 0.1 inch. Moreover, the vertical depth of the crescent-shaped deformation found on the 10-inch fracture of the airplane's aft bulkhead upper flange was exactly 0.1 inch. (NTSB 1979, 50) Tests performed by American Airlines confirmed that deformation of this magnitude would initiate an overload crack. This was strong evidence that the overload event that caused the 10-inch fracture of the upper flange indeed occurred during the maintenance operation on March 29-30, 1979.

- Final conclusion of the argument: The 10-inch overstress crack in the aft bulkhead's upper flange was created during a maintenance operation on May 29-30, 1979 at the American Airline's facility in Tulsa, Oklahoma, when the upper flange moved against the wing fitting clevis, and the clevis applied downward bending stress that deformed and cracked the upper flange surface.

7 Conclusion

The main goal of this paper is to show that (1) engineering failure investigations are a rich source for examining how to reconstruct the actual causes of past events, and (2) it is instructive to make explicit the structure of evidential reasoning used in the practice of engineering failure investigations. I have focused on a method called "feature dependence", examined the evidential basis of the key premises of feature dependence, and illustrated how a subtype of feature dependence called "descriptive enrichment" is used in practice.

I end this paper with two comments about the generalizability of my account of feature dependence. First, I have used the NTSB's investigations of aviation accidents as my main case studies of engineering failure investigations, and I believe that my account of feature dependence generalizes to other types of engineering failure investigations. I chose aviation examples because they tend to be well documented, and I chose the NTSB's investigations because the NTSB has an excellent reputation and is widely regarded as an authority in transportation safety (Fielding, Lo, and Yang 2010).

Second, outside of engineering failure investigations, historical scientists also encounter problems of reconstructing the actual causes of past events.[30] Even though I have described feature

[30] There is rich philosophical literature on the methodology of historical sciences, although much of the literature is focused on the methods of historical sciences in general and not specifically on reconstructing the actual causes of past events. As Turner 2009 and Currie 2019 point out, historical sciences are not limited to the studies of individual events; they also study regularities among event-types (e.g., processes of fossilization) and the causes of patterns in the traces (e.g., size and complexity increase in fossil records). For competing accounts of the methodology of

dependence as a method for reconstructing actual causes in engineering failure investigations, my account of the structure and evidential basis of feature dependence can be stated at a level of abstraction where specific engineering concepts and examples are abstracted away. This raises the question of whether my account of feature dependence applies to other types of research that also seek to reconstruct the actual causes of past events.

I believe that feature dependence plays a role in historical sciences, but I am not confident that my account of feature dependence based on engineering examples is fully applicable to historical sciences. There are major differences between epistemic situations of the engineering failure investigators and those of historical scientists, some of which could be relevant to how feature dependence works. For instance, one difference is that many historical scientists study events and processes of the *deep past*, whereas engineering failure investigators typically study events that occurred recently. Because historical scientists often study events and processes that happened a long time ago, many informative traces from these events and processes had already been erased by what Sober 1991 calls *information-destroying processes*, which reduces historical scientists' ability to access the features of past events. Information-destroying processes pose a less severe problem in engineering failure investigations, partly because the investigations occur right after the accidents, and partly because the engineering systems can be designed in ways that protect them from information-destroying processes to some extent during accidents.

A comparative study of how feature dependence works in engineering failure investigations versus in historical sciences is beyond the scope of this paper, but it could be a worthwhile pursuit for further work. Questions to be explored include: What is the relationship between features and traces? What types of feature dependence arguments are commonly used in historical sciences? Do the feature dependence statements used in historical sciences have a different evidential basis? How do feature necessity and feature sufficiency relate to existing accounts of the methodology of historical sciences? Answering these questions could deepen our understanding of the scope and limitation of feature dependence across the disciplines.

historical sciences, See Cleland 2002, Forber and Griffith 2011, Currie 2018, among others.

References

Brooks, Craig et al. (2003). "Predictive modeling of structure service life". In: *ASM Handbook* 13, pp. 946–958. DOI: `https://doi.org/10.31399/asm.hb.v13a.a0003706`.

Cleland, Carol (2002). "Methodological and epistemic differences between historical science and experimental science". In: *Philosophy of science* 69.3, pp. 447–451. DOI: `https://doi.org/10.1086/342455`.

Currie, Adrian (2018). *Rock, bone, and ruin: An optimist's guide to the historical sciences*. MIT Press.

— (2019). *Scientific knowledge and the deep past: history matters*. Cambridge University Press.

Fielding, Eric, Andrew W Lo, and Jian Helen Yang (2010). "The National Transportation Safety Board: A model for systemic risk management". In: *Available at SSRN 1695781*.

Forber, Patrick and Eric Griffith (2011). "Historical reconstruction: Gaining epistemic access to the deep past". In: *Philosophy and Theory in Biology* 3.201306, pp. 1–19. DOI: `https://doi.org/10.3998/ptb.6959004.0003.003`.

Galison, Peter (2000). "An accident of history". In: *Atmospheric flight in the twentieth century*. Springer, pp. 3–43.

Hall, Ned (2004). "Two concepts of causation". In: *Causation and counterfactuals*, pp. 225–276. DOI: `https://doi.org/10.7551/mitpress/1752.003.0010`.

Halpern, Joseph and Judea Pearl (2005). "Causes and explanations: A structural-model approach. Part I: Causes". In: *The British journal for the philosophy of science*, pp. 843–887.

Hanley, Brian (2021). "What caused the Bhopal Gas Tragedy? The philosophical importance of causal and pragmatic details". In: *Philosophy of Science* 88.4, pp. 616–637. DOI: `https://doi.org/10.1086/713902`.

Jacobson, Steven (2010). "Aircraft loss of control: Causal factors and mitigation challenges". In: *AIAA Guidance, navigation, and control conference*, p. 8007.

Lewis, David (1974). "Causation". In: *The Journal of Philosophy* 70.17, pp. 556–567.

— (2004). "Causation as influence". In: *The Journal of Philosophy* 97.4, pp. 182–197. DOI: `https://doi.org/10.2307/2678389`.

NTSB, National Transportation Safety Board (1979). *Aircraft Accident Report: American Airlines, Inc. DC-10-10, N110AA, ChicagoO'Hare International Airport Chicago, Illinois, May 25, 1979.* Tech. rep. Washington, DC.

— (1990). *Aircraft Accident Report: United Airlines Flight 232, McDonnell Douglas DC-10-10, Sioux Gateway Airport, Sioux City, Iowa, July 19, 1989.* Tech. rep. Washington, DC.

— (2001). *Aircraft Accident Report: Uncontrolled Collision With Terrain, United Airlines Flight 585, Boeing 737-291, N999UA, 4 Miles South of Colorado Springs Municipal Airport, Colorado Springs, Colorado, March 3 1991.* Tech. rep. Washington, DC.

Paul, L.A (2004). "Aspect causation". In: *The Journal of philosophy* 97.4, pp. 235–256. DOI: https://doi.org/10.7551/mitpress/1752.003.0009.

Paul, L.A and Ned Hall (2013). *Causation: A user's guide.* Oxford University Press.

Sober, Elliott (1991). *Reconstructing the past: Parsimony, evolution, and inference.* MIT press.

Turner, Derek (2009). "How much can we know about the causes of evolutionary trends?" In: *Biology & Philosophy* 24.3, p. 341. DOI: https://doi.org/10.1007/s10539-008-9139-5.

Woodward, James (2003). *Making things happen: A theory of causal explanation.* Oxford university press.

— (2010). "Causation in biology: stability, specificity, and the choice of levels of explanation". In: *Biology & Philosophy* 25.3, pp. 287–318. DOI: https://doi.org/10.1007/s10539-010-9200-z.