Concurrent Engagement and the Enclosed Application

Flyxion
Independent Researcher

November 2026

Abstract

Discussions of the Internet as a cognitive environment often move from the documented costs of task switching and interruption to the conclusion that digital environments as such fragment thought. This essay argues that evidence of switching costs underdetermines interface design: the same findings can motivate notification suppression, better support for resumption, or restrictions on the contexts a user may hold, and further premises are needed to choose among them. It distinguishes three practices grouped under the word concurrency (preserving inactive contexts, alternating between active tasks, and processing simultaneous media) and four properties that need not coincide: control, initiation, relevance, and resource competition. The studies reviewed here do not discriminate adequately among these arrangements. The essay also considers how circulating descriptions of attention may shape self-report. The design position defended combines evidence about attentional costs with a normative commitment to the user's authority over the purposes of their work: interfaces should make switching deliberate, preserve unfinished work, reduce unsolicited demands, and allow protected concentration. The essay does not claim that particular concurrent arrangements improve comprehension or performance.

1 Introduction

The standard complaint about the Internet as a cognitive environment has a recognizable structure. A reader who once followed an argument across many pages now scans, searches, and moves on, and the cause is located in the medium: links invite departure, alerts divide attention, and the habit of switching strengthens at the expense of the habit of staying. Carr gives this complaint its most influential form, and popular presentations of his book repeat it with additional emphasis, describing the page-scanning behavior of web readers and the cost of interruption as evidence of a general weakening of concentration [5, 1]. The target here is the complaint as it circulates in popular discussion, with Carr's book as its principal source, and not a full reconstruction of Carr's own argument. The argument is persuasive because it names something many readers have experienced. It is incomplete because the frameworks in which it is usually applied do not distinguish between kinds of switching, and because the same word, concurrency, is used for practices that differ in their cognitive demands.

A person working at a desktop computer under a modern windowing system, or in a browser that holds dozens of tabs across several sessions, occupies a different situation from a person whose phone issues a sound every few minutes at the behest of an advertising platform. In the first case the user usually decides what is open, what is in front, and what is running quietly behind. A text can be read in one window while a program is written in another, with a reference work in a third and a lecture or album available in a fourth. In the second case the structure of the environment has been written by someone else, and its purpose is to maintain a continuous claim on the person's attention. The difference between these situations is real, but it is not a single difference. It involves who controls the arrangement, who initiates each change of focus, whether that change serves the task in hand, and whether the streams involved draw on the same limited resources. These properties are analytically distinct and need not coincide, and treating any one of them as a proxy for the others produces mistakes in both directions. Concurrency names the general phenomenon, with the three specific practices distinguished in Section 3.

The thesis defended here concerns what follows from findings about switching costs. Those findings are well documented, but they underdetermine interface design: the same evidence can motivate notification suppression, better support for resumption, or restrictions on the contexts a user may hold, and further premises are needed to choose among them. The design position defended here combines evidence about attentional

costs with a normative commitment to the user's authority over the purposes of their work, together with the observation that platform incentives may conflict with those purposes. On that basis the essay supports interfaces that make switching deliberate, that preserve the state of unfinished work, that reduce unsolicited demands, and that allow periods of protected concentration. It makes no claim that specific concurrent arrangements, such as background lectures during programming, improve comprehension or performance, and it treats that as an open empirical question.

The essay proceeds as follows. Section 2 identifies the framework under examination and the distinction it omits. Sections 3 to 5 distinguish three senses of concurrency and engage the media multitasking literature, treat persistence and markers of encounter, and describe the architecture of user-controlled contexts with four analytically distinct properties. Sections 6 and 7 treat the enclosed application as a concentration of authority over attentional conditions and examine what the evidence discriminates. Sections 8 and 9 consider the vocabulary of attention as a looping classification and take up counterarguments, and Sections 10 and 11 draw design implications and conclude.

2 A Framework That Omits a Distinction

The argument under examination has three components that are usually presented together. The first is a claim about print: the book trains a capacity for prolonged, uninterrupted engagement with one line of thought, and that capacity supports interpretation and independent understanding [5]. The second is a claim about working memory. Because the capacity of working memory is limited, material that occupies it with navigation decisions, notifications, and parallel content leaves less available for comprehension and for transfer into long-term memory. The third is a claim about plasticity: since the brain changes with repeated activity, habitual scanning and switching strengthen those skills and leave the skills of sustained reading less practised.

Each claim has some support, and none by itself implies that the number of available contexts is a measure of fragmentation. The limitation of working memory does not require the view that every additional stream is a loss, and nothing in the argument as summarized here shows that Carr equates a dormant tab, instrumental music, a deliberate hyperlink decision, and an unsolicited notification. The object of analysis is the circulating form of the argument, instantiated here by one popular presentation. That presentation is a commentator's discussion and not Carr's own text. It is engaged here because it shows how the argument is received and extended in public discussion. It

provides evidence about that reception, not about Carr's theoretical commitments, and it is not offered as representative of every reader of the book. The popular presentation examined here does not consistently distinguish the availability of another context from the attentional demands of engaging it. The distinctions needed are developed in Sections 3 and 5, namely three senses of concurrency and four properties of an environment (control, initiation, relevance, and resource competition) that are analytically distinct and need not coincide. A paper, a terminal, a reference work, and a paused lecture can coexist on a screen without all of them demanding active processing, and a framework that counts open items as fragmentation will misdescribe that situation.

The popular presentation shows the difficulty in its own terms. It describes tabs accumulating until the original purpose is forgotten, treats the scanning of web pages as the normal condition of online reading, and warns against a hyperdistracted style of use, while also appearing to recommend, in its speaker's own practice, short bursts of intense focus, movement among several books, and the use of the Internet as a research tool for those who know what to look for [1]. The speaker's habits are therefore a mixture of the nonlinear and the concentrated, which an indiscriminate view does not predict. The more careful reading of that discussion is that movement among sources is compatible with understanding when the reader retains a question and returns to unfinished problems, and obstructive when successive diversions repeatedly displace the purpose that made the material relevant. That is a claim about the relation between a reader and a purpose, and it cannot be settled by counting streams.

A further assumption lies behind the contrast with the book. The book against which the web is measured is the book read from cover to cover by a person with nothing else open. Much of the history of the codex is a history of devices for not reading in that way, and textbooks, dictionaries, manuals, and collections of essays are routinely entered at the middle, consulted in fragments, and set beside other books on the same desk. The apparatus of tables of contents, alphabetical order, indexes, and concordances, much of it developed from the twelfth century onward and elaborated in the early modern period, was built so that a reader could find a passage rapidly without reading what precedes it [17, 18]. This does not make linear reading unimportant. It shows that the geometry of a text, linear or nonlinear, is not the variable that separates depth from shallowness.

Rapid navigation can itself serve an intellectual purpose. A reader moving through an encyclopedia by following previously unvisited links may be exploring the organization of a domain, encountering names, distinguishing categories, and answering successive local questions. A brief visit may suffice to read an opening definition or identify a

relation without providing an understanding of the article as a whole, and comprehending the whole article was not necessarily the task. Where a simple navigation rule guides the sequence, such as following the first link not yet visited, continuity can reside in the exploratory procedure even as the subject changes, and the record of visited links preserves part of the route. Such activity should be assessed by what it enables the reader to identify, connect, or subsequently investigate, and not by dwell time alone. This does not establish durable learning from brief exposure. It distinguishes an exploratory reading task from an unsuccessful attempt to read each article comprehensively, and it means that a governing task may be a question that changes as the reading supplies information.

3 Three Senses of Concurrency

The word concurrency is used here, and in the surrounding discussion, for three practices that need to be distinguished because they draw on different evidence and carry different demands. The first is preserving inactive contexts for later return. A reference tab, a half-read paper, and a terminal session left open while another task proceeds are preserved without necessarily being processed. Preserving a context is a matter of external organization, and processing it is a matter of attention; the two are separable in principle, even where they interact in practice, as when a visible title prompts recollection of an unfinished task. This concerns external organization and the preservation of work, and its cognitive significance is mainly that it reduces the burden of reconstructing where one was. The second is alternating between active tasks, in which attention moves from one engaged activity to another and back. Here the relevant literature is that of switching and resumption, which documents costs and also shows that they depend on factors such as preparation time and the similarity of the tasks [12]. The third is processing simultaneous media, in which two or more streams are attended to, even peripherally, at the same time. The relevant literature here concerns interference between concurrent demands.

Evidence against the third practice does not automatically count against the first. A finding that heavy simultaneous media consumption is associated with poorer performance on certain attentional measures says little about whether leaving reference tabs open helps or hinders a writer. The converse also holds. The usefulness of persistent tabs does not establish that a familiar film or lecture helps someone program or read. Each practice has to be assessed on its own terms, and the essay’s positive claim concerns chiefly the

first, with the second as a matter of managing costs and the third as unsettled.

The media multitasking literature is the strongest source of caution. Ophir, Nass, and Wagner reported that heavy media multitaskers performed worse than light multitaskers on several tasks involving the filtering of irrelevant information and switching between tasks [15]. Later reviews have found that the associations are heterogeneous across studies and measures, that the direction of causation is unresolved, and that results depend on what is multitasked and how it is measured [21, 22]. Two points follow. First, much of this research concerns habitual simultaneous consumption of competing content, such as messaging while reading, and this places two streams in the same linguistic and executive channel. Accompaniments differ in their overlap with the demands of the focal task, and their effects cannot be inferred from their mere presence. Instrumental music, a lecture on a familiar subject, and a film that has been seen several times might impose very different loads, and low intended attention or familiarity does not establish an absence of processing demands. The irrelevant speech effect, in which unattended speech disrupts verbal short-term memory, is one reason to suspect that verbal accompaniment is not free [19]. Second, user control permits experimentation, but it does not ensure that users discover which arrangements improve comprehension or performance. They may discover what feels comfortable, what makes work more tolerable, or what sustains their participation in a long task. Each of these can matter, since a task that is tolerable is more likely to be continued, but none is identical to improved accuracy or understanding. For this reason the question whether particular concurrent arrangements help is treated here as open.

4 Persistence and Markers of Encounter

The first sense of concurrency, preserving inactive contexts, deserves fuller treatment than the other two, since it has a dimension that a focus on simultaneity obscures. A bookmark, a written page number, or a recorded timestamp allows an engagement to remain unfinished without being abandoned. What persists is not continuous attention but a recoverable relationship to the material, and returning may require some rereading or listening again, a cost that differs from reconstructing the source, locating the passage, and remembering why it mattered. Two things are worth preserving: the position in the source and the state of the inquiry. A bookmark records where a reader stopped, while a note such as a remark that a passage may explain a discrepancy in an experiment records what was being pursued. Neither guarantees recovery of the earlier understanding, but

together they provide an external scaffold for resumption, and digital persistence extends practices already available with physical books and does not introduce an entirely different kind of intellectual activity. The distinction also shows what restoring a workspace can and cannot do. An open document can retain its cursor position without recording what its reader intended to do next, so persistence of state is not yet persistence of intention.

The temporal scale of sustained inquiry also matters here. Sustained inquiry cannot generally mean uninterrupted concentration over years or decades. It means repeatedly recovering a problem, incorporating intervening experience, revising earlier judgments, and carrying unresolved questions forward, and a person may work on something else for a long time and then recognize a connection that makes an old passage newly intelligible. The continuity belongs to the developing inquiry across episodes of attention, and not to any single episode. This does not require the claim that switching is cognitively free. It indicates that local concentration and long-term intellectual continuity are related achievements that are not interchangeable, and that an environment should be assessed by whether it helps people preserve sources, questions, annotations, and unfinished intentions across departures, as well as by how it supports focus within a session.

Persistence also takes the form of markers of prior encounter, which work even when the marked item is never revisited. A bookmark that appears among the suggestions as a query is typed indicates that the item was found before and makes its recurrence more likely. A starred or watched article on a reference site records that the page has been attended to. A followed account or a forked repository can serve as a marker of prior encounter. A recorded commit identifier additionally identifies the repository state encountered and provides a reference against which later changes can be compared. Platform features that resurface items on anniversaries can serve a similar function, including for entries visible to no one else, which can then function as intervals of re-exposure. Research on distributed practice shows that spacing re-exposure supports verbal retention, with the effective interval depending on how long the material must be retained [6]. That research does not test the extension to unfinished projects, which is proposed here as an application and not as a result. None of these markers requires that the item be read again, and each lowers the cost of recognizing it as already known. The last example is instructive for the argument about enclosure. A feature built to retain attention can serve the owner's purposes when the owner controls what is stored and when it returns, which supports the claim that control, relevance, and initiation, and not enclosure as such, determine whether a mechanism helps or intrudes.

5 User-Controlled Contexts and Four Distinct Properties

What distinguishes the desktop and the open browser from earlier media is that they provide persistent parallel contexts that the user can arrange. A windowing system on Linux or Windows maintains several applications in distinct regions of the display, each with its own state, and permits the user to bring any of them to the foreground in a single action. Virtual desktops group windows into workspaces corresponding to different tasks. Multiple user sessions, terminals, and remote connections allow one machine to host several independent working environments. A browser with tabs and separate profiles reproduces this structure within a single application, with the added property that tabs can preserve resources from different servers and institutions, so that a reference work, a repository, a discussion forum, and a streaming video can be held open together.

Three features of this architecture matter. The first is persistence: a context that has been left is not destroyed, and the user can return to it where it was. The second is that movement among contexts is, with exceptions, caused by a decision of the person. The third is graded visibility: a window can be full-screen, tiled at the edge, minimized, or running in the background with only its audio audible. These features constitute a toolkit for organizing work, and they are the basis of the claim that the ability to preserve and revisit contexts has value independent of any particular cognitive benefit.

Several influential traditions in computing developed resources for user-directed retrieval, association, and the preservation of multiple working contexts, and this lineage is relevant to what a design should protect, although it does not establish a unified original purpose. Interactive time-sharing was motivated by letting many people use a machine at once, which is not the same as supporting the concurrency of one person’s activities, and early statements of human-computer interaction described a partnership in which the machine would handle routine retrieval and bookkeeping so that the person could concentrate on formulation and judgment [10]. The hypertext tradition, from Bush’s proposal for a personal device of associative trails to Nelson’s account of linked, changing documents, treated the reading of connected material as a primary use [4, 13]. Associative navigation of this kind is not identical to simultaneous engagement, since a trail is followed in sequence. The Web then made any document on any server reachable from any other through a link, and the tab later made it possible to hold several such documents open at once.

The lineage is heterogeneous, and the fit of individual sources is partial, since

Licklider's account concerns routine retrieval and bookkeeping more than the arrangement of several contexts. This history supports the design possibilities defended here, since it shows that retrieval, association, and the preservation of several working contexts were treated as features of capable systems, and that they became part of ordinary desktop and browser practice. It is offered as background and not as a claim about an original purpose of the network.

The argument requires that four properties be kept apart, because the literature and the popular discussion often move among them without notice. The first is control over the environment: who determines which contexts exist, how they are arranged, and what is permitted to appear. The second is initiation of a particular switch: whether the person chose to move or the system moved them. The third is relevance to the governing task: whether the content that draws attention bears on what the person is trying to do. The governing task may itself be exploratory, as discussed in Section 2, so that relevance can be assessed against a question that changes as the reading supplies information. The fourth is competition for cognitive resources: whether the streams involved overlap in the processing they require, for example by both demanding verbal comprehension. These are conceptually separable, although actual systems may couple them, and none reliably substitutes for the others. An open desktop can deliver incessant notifications, and an enclosed application can provide uninterrupted reading. A user can voluntarily enter a feed and become distracted, while an externally initiated interruption can be directly relevant to an ongoing purpose, as when a colleague's message answers the question one was working on. In an experiment involving an email task, interrupted participants completed the focal work faster, excluding time spent on interruptions, but reported greater stress, frustration, time pressure, and effort. The experiment compared same-topic and different-topic interruptions and found no significant difference between those conditions in task-completion time [11]. Shared topic, however, does not necessarily mean that an interruption advances the governing task. An activation-based model of memory for goals attributes difficulty in resuming an interrupted task to interference and to the availability of cues for pending goals [2], and a separate literature on task switching documents costs when the task changes [12], including costs when the change is chosen by the participant [3]. Neither body of work implies that every combination of the four properties is equally costly, and neither implies that user initiation removes cost. It is a separate question, which the studies seldom isolate, how much each property contributes.

Practitioner descriptions of terminal work show that user-controlled contexts are

layered, and that each layer has its own persistence. At the level of the operating system, a keyboard shortcut such as Alt and Tab moves among windows. Within an application, tabs and panes divide a single window, and one recent account of Windows Terminal describes colour-coded tabs, split panes opened from profiles, and the reopening of the previous session's tabs and panes at launch, including remote sessions [20]. Below the application, terminal multiplexers such as tmux, and front ends such as Byobu that wrap it, keep sessions running on a machine independently of any particular window or connection, so that a person can detach, disconnect, and later reattach to processes that continued in the interim, while the host and multiplexer server remain running. The layers differ in what survives a departure. A window switch preserves whatever the application holds, a restored session preserves the arrangement of tabs, and a multiplexer preserves running processes. A person who combines them chooses, at each layer, how much state to keep and how easily to return to it, which is the kind of control over arrangement that the four properties distinguish. The same account recommends the multiplexer for persistence on remote machines, which indicates that practitioners treat local session restoration as insufficient for some work. This is evidence about the infrastructure that practitioners build and use, and not evidence about cognitive effects.

A further distinction separates concurrency among machines from concurrency in attention. A person can keep sessions open on many separate computers, but the number of such contexts is bounded by hardware that fails, needs power, and must be maintained. Containers and virtual machines allow distinct system environments to share a host, reducing the need for a separate physical machine for each environment while retaining limits imposed by memory, processing capacity, storage, and maintenance. Multiplexers can also broadcast input to several panes, allowing the same command to be submitted to multiple sessions. These facilities increase the number of environments that can be maintained and operated without establishing that their contents are simultaneously processed by the person. Broadcasting shares an input across sessions, while interpreting their potentially different results still requires attention. The example indicates that the cost of preserving and operating many contexts has fallen, and that a remaining constraint lies in how much of that activity a person can monitor and interpret. It follows that counting open contexts or running processes, as a proxy for the load on attention, would confuse what the machines are doing with what the person is doing.

6 The Enclosed Application as Concentrated Authority

The later development that this essay treats as the central structural concern is the enclosure of the user within applications that are built to retain them. Social platforms and the mobile applications that followed them present a feed of content selected for its capacity to hold attention, discourage the departures that open links made easy, and put the feed, which is not stably addressable or archivable as a whole, in the position once held by the page, even though individual posts can often be linked and cited. The commercial logic has been described at length: where revenue depends on advertising exposure, platforms have incentives to prolong engagement, although time spent is only one determinant of commercial value [23]. That account concerns chiefly advertising-supported media, and subscription and commerce models create different incentives that would need separate analysis.

It is important to state what enclosure is and is not. It is not itself a cognitive mechanism. It is a political and architectural property, namely the concentration in the platform of authority over the conditions distinguished above: what appears, when it appears, how easily one can leave, and whether content can be moved to another context. Zittrain's distinction between generative systems, which accept uses their designers did not anticipate, and appliance-like systems, which are fixed in function and controlled by their vendor, describes the structural change [24]. A generative personal computer lets its owner install anything, open anything, and arrange the pieces as they wish, while an appliance offers a defined set of actions. The enclosed application tends toward the appliance. The consequence for attention is indirect. Because the platform controls the arrangement, it can choose to interrupt, to resist departure, and to select content for engagement, and it has commercial reasons to do so. Whether it does depends on its incentives and design, and an enclosed environment built for reading, with no notifications, would raise different questions. The claim is not that enclosure causes distraction, but that it limits the user's ability to alter those conditions while reserving important decisions to the provider, a party whose interests may conflict with the user's. Providers may also have interests aligned with users, such as reputation and subscription retention, so the claim concerns the possibility and structure of conflict and not its inevitability.

The contrast can be illustrated by four architectural features, which are examples of how control is implemented and do not replace the four properties distinguished in Section

5: where state is stored, how it is indexed, how predictable return is, and who governs injection and pacing. In the desktop and browser environments of the 1990s, which relied on multiple windows and on bookmarks and history kept as local files, with tabbed browsing becoming widespread only in the following decade, state was typically held on the user's machine, indexed by spatial and hierarchical arrangement such as windows, folders, and saved lists, and returned to with reasonable predictability, while the user governed what opened and when. Feeds tend to differ along each dimension. State is held largely on the provider's servers, indexed by recency or ranking, return to a prior position is less predictable because content is reordered and refreshed, and injection and pacing are governed by the platform. These are tendencies and not guarantees. Pages reload, browsers discard tabs, sessions fail to restore, and many open-web resources depend on remote servers, while feeds often offer saved items, drafts, and history, and a desktop crowded with notifications can impose the interruptions described above. The comparison therefore concerns defaults and who holds authority over them. The suggestion that particular arrangements may lower the cost of resuming a goal is a reasoned application of the memory-for-goals model and remains conjectural in this setting.

7 What the Evidence Discriminates

A reassignment of the evidence is tempting, but the sources do not support as much as it might appear. The best-known description of how people read on the web, the F-shaped scanning pattern, derives from eye-tracking of users encountering text-heavy pages, with the observation that attention concentrates on the opening lines and declines down the page [14]. It describes one common pattern on a particular genre of page under particular conditions, usually when people are looking for something, and the originating research group subsequently clarified that the F-pattern is one of several reading patterns and that its occurrence depends on page presentation and the reader's motivation and objectives [16], so it is used here only as an illustration of task-dependent scanning. It says little about what the same people do when they have chosen to read a long text. It therefore weakens an indiscriminate condemnation, because a scanning strategy suited to a task is not a deficit, and it does not support any positive causal account of enclosure.

The research on screens and comprehension supports a modest conclusion. A meta-analysis comparing paper and screen found a small overall advantage for paper in reading comprehension. The advantage was greater under time constraints and was found in studies using informational or mixed texts, but not in studies using narrative texts alone

[8]. This is relevant but limited. It concerns the effect of a medium on a single reading task and does not address purposeful movement between a text and a reference work. For the same reason it does not identify enclosure as the explanation of anything. The media multitasking research concerns associations whose causal interpretation remains unresolved [21], and the interruption research shows costs in stress [11] and offers a theoretical account of difficulty in resuming a task [2], in settings that often impose the interruption. Taken together, these findings can undermine a view that digital environments are shallow as such, and they can support caution about specific practices. Some studies isolate particular distinctions: Mark and colleagues, for example, manipulate the topical relationship between an interruption and the focal task [11]. The studies reviewed here do not, however, jointly isolate environmental control, initiation, relevance to the governing task, and resource competition across the arrangements distinguished in this essay, and targeted comparisons isolating particular relationships among the four properties would be needed to say more.

The plasticity argument can be treated similarly. If the brain adapts to the practices it performs most, then what it adapts to depends on the environment and on the person's use of it. A user who spends most of their time in a feed that delivers a rapid sequence of unrelated items practises something different from one who moves between a text, a reference work, and a program. This is a plausible reading, but it is a hypothesis, and the general claim that heavy Internet use weakens concentration is not established by the studies reviewed here.

8 The Vocabulary of Attention as a Looping Classification

Descriptions of attention do not only report experience; they also circulate, and in circulating they alter what is reported. Hacking's account of looping effects names the mechanism: a classification of people enters public use, those classified or tempted to classify themselves interpret their experience through it, and their altered behavior and self-description in turn change what the classification is taken to mean [9]. Whether a particular circulating description has produced such an effect is an empirical question that the concept does not settle, and the reinterpretation of object permanence in online discussion has been analyzed as a looping effect on the basis of interview evidence that its author describes as preliminary [7]. The popular presentation examined in this essay

supplies examples of the vocabulary involved, since it describes a brain as "fried," a reader as "ADHD brained," and a scattered style of attention as "hyper-dottiness," and phrases of this kind travel well because they are recognizable and carry emotional force. The speaker uses these terms descriptively and not as diagnostic claims. The point here concerns how such terms circulate and not whether any diagnostic category is valid, a question this essay does not address.

Carr's thesis can be treated in the same way as a description that enters circulation. A reader who has felt distracted encounters an account on which the Internet is reshaping attention, recognizes the experience, and adopts the account as an interpretation, and that uptake may encourage further descriptions of the same kind. This is a hypothesis about the reception of the argument and not a demonstrated process, and it does not show that the underlying claims are false. It does imply a methodological caution. If the available vocabulary shapes what people report, then self-reports of attentional failure are imperfect evidence about attention, since they reflect the categories through which experience is described as well as the experience.

Circulating vocabulary carries a second hazard, conflation, which differs from looping: a term acquires a meaning that merges two distinct claims. A concrete case concerns the term object permanence. In developmental psychology it names the understanding that an object continues to exist when it cannot be perceived, and in some online usage it has come to describe forgetting about things that are out of sight. These are different claims. Failing to remember a thing is not a belief that it has ceased to exist, and the confusion is instructive for the present argument, because the remedies differ. A person who forgets an open tab need not have lost the understanding that it continues to exist. The difficulty may instead concern retrieving its location, its relevance, or an unfinished intention. A note, bookmark, or recorded next action can preserve information useful for recovering that relationship, although its effectiveness depends on the circumstances of return. A vocabulary that merges the two directs attention toward the person and away from the environment's capacity to preserve context. The same lesson applies to counts of open contexts, since a single descriptor conflates conditions that the four properties in Section 5 keep apart. Only conflation bears directly on the design argument, since looping concerns how reports are shaped. The caution applies reflexively as well: descriptions of preserved contexts and unfinished intentions offered here are proposed as analytic categories and not as reports, and they would need the same scrutiny if they were used as evidence.

The same case shows why the conditions of forgetting matter. A report that someone

cannot remember what they were doing is difficult to interpret without information about the task, intervening activity, elapsed time, and available retrieval cues. An activation-based model of memory for goals explains retrieval in terms of activation, interference, and associative priming [2]. The number of open contexts is therefore a potentially relevant descriptor of the environment, but it is not itself a measure of cognitive load: preserved contexts may remain inactive, while a single competing activity may impose substantial demands. Research should therefore record the variables specified in Section 10, which bear on the four properties identified in Section 5, although causal comparisons would require further experimental control. Discussions of the phenomenon should likewise distinguish a failure to retain a goal from a failure of the environment to persist the context in which the goal was held.

9 Counterarguments

Several objections deserve a direct response. The first is that even peripheral media costs something. This is probably true. Background audio or video with a verbal component can compete with reading, and a person's belief that they work better with such accompaniment is not reliable evidence that they do. The response is not that the practice is free but that its costs are unsettled and depend on overlap with the task, which is why the essay does not claim a benefit. The second objection is that most users lack the self-regulation that open systems presume, which would put the design position in tension with itself: it defends user control while conceding that users may exercise it poorly. The tension is real and should be stated plainly. The response has two parts. First, the position does not defend unstructured freedom. It asks for defaults and affordances that support the person's own purposes, such as protected workspaces, deliberate switching, and notifications that are off unless requested, and these are forms of support that leave the person in charge of what the purposes are. Second, if self-regulation is scarce, the conclusion is not that authority over attention should pass to a party with commercial interests in it, since that party's interests may conflict with the user's goals. It is that control should remain with the user while the interface lowers the cost of exercising it well. The remaining difficulty is that people differ, and a default suited to one user may constrain another, which is a reason to make such supports adjustable and not to build them in as fixed policy. The third objection is that enclosure responds to user demand and that people prefer the feed. Preference revealed under conditions engineered to produce engagement is a weak guide to what people would choose elsewhere, but it is not

nothing, and not all use of enclosed applications is imposed. The same skepticism applies to claims that users want protected concentration, which this essay does not assume.

A further objection is that the position is paternalistic, since offering protected concentration assumes that users need help managing their own attention. The reply is that the supports proposed are transparent, adjustable, and reversible at low cost, so that the person decides whether and when to use them, and that the alternative to such supports is not neutrality, since existing environments already embody someone's decisions about attention. Every functioning interface must establish some initial configuration, so the relevant question is whether an arrangement overrides a person's choices for their supposed benefit, and defaults that the user can see and change at low cost do not. Reversibility is itself a design requirement and not a guarantee, since a setting buried in menus is adjustable only in name. Reversible defaults are the design answer to the tension between defaults that matter and settings that differ by person, because they relocate authority to the user without requiring the user to author the initial configuration. A related worry is that the principles ask a great deal of designers, whose incentives often point the other way. That is a reason to ask where authority over these defaults should sit, and it is one motive for preferring open, generative systems in which the user or a third party can supply the supports that a vendor declines to provide.

10 Design Implications

The principles below combine three kinds of premise. Evidence on switching and interruption identifies costs that design may seek to reduce. The commitment to the user's authority over the purposes of their work guides the choice among possible responses, including deliberate switching, adjustable settings, and addressable and exportable content. Further proposals, notably that external records of position and intention ease resumption in these environments, are reasoned applications of models of memory for goals and not results established in this setting.

If switching costs are real, and if the user is to retain authority over the purposes of their work, several design principles follow. Interfaces should make switching deliberate, so that a change of context is an action the person takes and can see, not a side effect of an alert. They should preserve the state of unfinished work, so that leaving a task costs little and return costs less, since models of memory for goals emphasize cues for retrieving pending goals [2], which suggests, as a reasoned application and not a tested result, that an external record of where one was reduces the burden of reconstructing

them. Because restoring a workspace does not restore an intention, they should also offer optional cues about the unfinished task, such as a brief note, a marked passage, or a recorded next action. They should reduce unsolicited demands, with the default that nothing interrupts unless the person has asked it to, and notifications governed by the user's declared purposes and not by the platform's interest in recall. They should allow protected concentration, in the form of modes or workspaces in which a chosen set of contexts is available and everything else is held back. Content should be addressable, so that a link can lead to it from anywhere and a reader can return to it, and applications should not penalize departure or prevent content from being exported, saved, or opened beside other material. Finally, operating systems and browsers should continue to support the generative arrangement that makes persistence and arrangement possible, including multiple windows, sessions, and profiles. Research and evaluation of such interfaces should record the number and visibility of preserved contexts, available resumption markers, the user's control over the arrangement, who initiated each switch, the relevance of the switched-to material to the governing task, and the processing demands of concurrently engaged activities, so that the four properties can be distinguished, although causal comparisons would require further experimental control. These measures aim at an environment that leaves the allocation of attention with the person who has to live with it, while supporting that person's efforts to concentrate.

11 Conclusion

Carr's diagnosis of shallowness names a real concern, and the evidence on switching and interruption gives it a genuine basis. What the evidence does not do is determine the design response. It is compatible with suppressing notifications, supporting resumption, or restricting the contexts a user may hold, and the choice among these rests on further premises, including a commitment to the user's authority over the purposes of their work. The open desktop and the tabbed browser continue an older practice of keeping several books open and consulting indexes, dictionaries, and cross-references, and they add persistence and arrangement under the user's control. Assessing threats to sustained thought requires distinguishing control, initiation, relevance, and resource competition, rather than treating the medium as a sufficient explanation, and the enclosed application matters because it concentrates authority over those properties in a party with its own interests. The design position set out in Section 10 combines that evidence with that commitment. The question that remains for designers is therefore not whether the

Internet makes people shallow, but whether environments are built on the assumption that attention is a resource to be captured or on the assumption that it is a capacity to be supported.

References

[1] Academic Agent. *The Internet Is Making Us Stupid, Nicholas Carr*. Video lecture, YouTube, 14 October 2026. `https://youtu.be/wEozfS-11WM`

[2] Altmann, E. M. and Trafton, J. G. Memory for Goals: An Activation-Based Model. *Cognitive Science*, 26(1):39–83, 2002.

[3] Arrington, C. M. and Logan, G. D. The Cost of a Voluntary Task Switch. *Psychological Science*, 15(9):610–615, 2004.

[4] Bush, V. As We May Think. *The Atlantic Monthly*, 176(1):101–108, 1945.

[5] Carr, N. *The Shallows: What the Internet Is Doing to Our Brains*. W. W. Norton, New York, 2010.

[6] Cepeda, N. J., Pashler, H., Vul, E., Wixted, J. T., and Rohrer, D. Distributed Practice in Verbal Recall Tasks: A Review and Quantitative Synthesis. *Psychological Bulletin*, 132(3):354–380, 2006.

[7] Chevalier, O. "It Starts on TikTok": Looping Effects and the Impact of Social Media on Psychiatric Terms. *Philosophy, Psychiatry, & Psychology*, 31(2):163–174, 2024.

[8] Delgado, P., Vargas, C., Ackerman, R., and Salmerón, L. Don't Throw Away Your Printed Books: A Meta-Analysis on the Effects of Reading Media on Reading Comprehension. *Educational Research Review*, 25:23–38, 2018.

[9] Hacking, I. The Looping Effects of Human Kinds. In D. Sperber, D. Premack, and A. J. Premack, editors, *Causal Cognition: A Multidisciplinary Debate*, pages 351–383. Clarendon Press, Oxford, 1995.

[10] Licklider, J. C. R. Man-Computer Symbiosis. *IRE Transactions on Human Factors in Electronics*, HFE-1(1):4–11, 1960.

[11] Mark, G., Gudith, D., and Klocke, U. The Cost of Interrupted Work: More Speed and Stress. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 107–110. ACM, 2008.

[12] Monsell, S. Task Switching. *Trends in Cognitive Sciences*, 7(3):134–140, 2003.

[13] Nelson, T. H. Complex Information Processing: A File Structure for the Complex, the Changing and the Indeterminate. In *Proceedings of the 1965 20th National Conference*, pages 84–100. ACM, New York, 1965.

[14] Nielsen, J. F-Shaped Pattern for Reading Web Content. Nielsen Norman Group, Alertbox, 16 April 2006.

[15] Ophir, E., Nass, C., and Wagner, A. D. Cognitive Control in Media Multitaskers. *Proceedings of the National Academy of Sciences*, 106(37):15583–15587, 2009.

[16] Pernice, K. F-Shaped Pattern of Reading: Misunderstood, but Still Relevant (Even on Mobile). Nielsen Norman Group, 12 November 2017. `https://www.nngroup.com/articles/f-shaped-pattern-reading-web-content/`

[17] Rouse, R. H. and Rouse, M. A. Statim Invenire: Schools, Preachers, and New Attitudes to the Page. In R. L. Benson and G. Constable, with C. D. Lanham, editors, *Renaissance and Renewal in the Twelfth Century*, pages 201–223. University of Toronto Press, Toronto, 1991.

[18] Saenger, P. *Space Between Words: The Origins of Silent Reading*. Stanford University Press, Stanford, 1997.

[19] Salame, P. and Baddeley, A. Disruption of Short-Term Memory by Unattended Speech: Implications for the Structure of Working Memory. *Journal of Verbal Learning and Verbal Behavior*, 21(2):150–164, 1982.

[20] Sherback, T. I Made Windows Terminal My Home for PowerShell, WSL, and SSH and Stopped Juggling Three Windows. XDA Developers, 3 November 2026. `https://www.xda-developers.com/i-made-windows-terminal-my-home-for-powershell-wsl-and-ssh-and-stopped-juggling-three-windows/`

[21] Uncapher, M. R. and Wagner, A. D. Minds and Brains of Media Multitaskers: Current Findings and Future Directions. *Proceedings of the National Academy of Sciences*, 115(40):9889–9896, 2018.

[22] van der Schuur, W. A., Baumgartner, S. E., Sumter, S. R., and Valkenburg, P. M. The Consequences of Media Multitasking for Youth: A Review. *Computers in Human Behavior*, 53:204–215, 2015.

[23] Wu, T. *The Attention Merchants: The Epic Scramble to Get Inside Our Heads.* Alfred A. Knopf, New York, 2016.

[24] Zittrain, J. *The Future of the Internet and How to Stop It.* Yale University Press, New Haven, 2008.