Life Events as a basis for PKG Agent Friendly Registries

Aisling Third[1], John Domingue[1]

[1]Knowledge Media Institute, The Open University, Walton Hall, Milton Keynes, MK7 6AA, UK

Abstract

Personal Knowledge Graphs (PKGs) are increasingly used to support user-centric AI systems, yet current approaches provide limited guidance on how autonomous agents can efficiently navigate, interpret, and protect the heterogeneous personal data they contain. This paper proposes a psychologically and sociologically grounded Life Events Registry Ontology as a lightweight top-level semantic layer that enables trustworthy agent-PKG interaction. Drawing on established research in autobiographical memory, cultural life scripts, and life course sociology, we argue that life events constitute cognitively natural, socially recognisable, and institutionally meaningful structures for organising personal information. We operationalise this insight by defining a minimal set of LifeEvent primitives—spanning lifetime periods, general events, and event-specific detail—supported by contextual anchors and personal goals. This abstraction enables a computer-using personal AI agent (CUPA, [1]) to perform rich private inference locally while exposing only high-level, privacy-preserving signals externally. We illustrate how the model extends across major life domains (e.g. health and education) and show how it supports data minimisation, interpretability, and institutionally aligned agent behaviour. The resulting framework provides a cognitively plausible, operationally lightweight, and privacy-by-design foundation for developing trustworthy PKG-based personalised agents.

Keywords

Personal Knowledge Graphs, Agents, Education, Privacy by Design, Life Events

1. Introduction

1.1. Personal Knowledge Graphs and Agent Interaction

Personal Knowledge Graphs (PKGs) have emerged as a promising paradigm for representing heterogeneous, user centric data in a structured and machine interpretable form. PKGs aim to integrate information spanning identity, activities, relationships, preferences, and personal records, often under conditions of user control and data sovereignty. This vision is increasingly realised through decentralised personal data stores, such as Solid Pods, where individuals retain ownership over their data while enabling selective access by applications and agents [2]. Recent advances in AI have intensified interest in agent based interaction with PKGs, where autonomous or semi autonomous agents act on behalf of individuals to support tasks such as planning, recommendation, compliance checking, or personal assistance. However, while existing PKG research has focused extensively on data modelling, interoperability, and privacy, comparatively little attention has been paid to the entry level semantics required for effective agent navigation. In practice, agents are often expected to query large and heterogeneous graphs without a shared understanding of where relevant information is likely to reside or which subsets of the PKG are appropriate for a given task. As a result, agent interaction with PKGs frequently relies on application specific assumptions, hard coded schema knowledge, or overly broad queries that increase both computational overhead and privacy risk. This highlights the need for an intermediate semantic layer that supports agent interpretable discovery, scoping, and access control, without imposing a rigid or overly abstract upper ontology.

ESWC'26: Trust, Autonomy and Accountability in PKG-Based Agentic AI (TAAPAAI) Workshop on May 10, 2026 in Dubrovnik, Croatia

aisling.third@open.ac.uk (A. Third); john.domingue@open.ac.uk (J. Domingue)

0000-0002-0386-1936 (A. Third); 0000-0001-8439-0293 (J. Domingue)

1.2. Life Events and the Structure of Autobiographical Knowledge

Insights from cognitive psychology and memory research suggest a compelling foundation for such an intermediate layer. Research on autobiographical memory consistently shows that human recollection and self understanding are organised around life events and culturally shared expectations about the life course. Berntsen and Rubin's influential work on cultural life scripts demonstrates that individuals across a culture share relatively stable expectations about major life events—such as education, family formation, career transitions, health events, and relocation—and that these expectations strongly shape how autobiographical memories are encoded and retrieved [3]. More broadly, models of autobiographical memory describe a hierarchical structure in which life periods, general events, and event specific knowledge form the backbone of personal knowledge organisation [4]. Importantly, these structures are neither arbitrary nor purely individual; they reflect socially and institutionally grounded domains of life, including education systems, healthcare, employment, family relations, and belief systems. Subsequent work has shown that these life structures display individual variation, allowing for both shared semantic anchors and personal specificity [5]. From the perspective of PKGs, this literature suggests that life events provide a cognitively plausible and socially grounded indexing mechanism for personal data. Rather than treating personal information as an undifferentiated collection of triples, life event framing offers a way to align PKG organisation with how individuals—and, by extension, agents acting for them—conceptually segment their lives.

1.3. Life Course Theory and Institutional Life Domains

Complementary insights arise from sociology, particularly life course theory, which studies how individual lives unfold through sequences of socially structured roles, transitions, and institutions. Life course scholars emphasise that lives are shaped by participation in key societal domains—such as education, work, family, health, and religion—and by transitions between roles within these domains [6, 7]). The concept of linked lives further highlights that relationships and family are central organising forces across the life course.[8]

Crucially, life course research frames these domains as institutionally meaningful units, governed by norms, policies, and expectations. Education, healthcare, finance, and family relations are not merely thematic groupings; they correspond to distinct institutional contexts with well defined boundaries and obligations [9]. This perspective is particularly relevant for agent mediated PKGs, as many agent tasks—such as eligibility checking or compliance support—implicitly operate within these institutional domains. Adopting life course informed categories as top level PKG entry points therefore aligns agent reasoning with the same structural divisions that organise social systems and public services.

1.4. Life Events as Registries in Digital Systems

The use of life events as organisational primitives is not limited to cognitive or sociological theory. In e-government, life event based models have been adopted to structure citizen services around major transitions such as birth, marriage, relocation, employment, and retirement. Prior work has shown that life events function effectively as service registries, aggregating heterogeneous data and processes around user centric milestones rather than administrative silos [10, 11].

This approach has been promoted at scale in European e-government initiatives, where life events are used to guide service discovery and data integration across agencies. Notably, these systems treat life events not as domain ontologies, but as navigation and coordination mechanisms sitting above specific models.

This observation motivates our proposal to generalise the life event approach from public sector service delivery to private, agent mediated PKGs. In this context, a life events based top-level ontology serves as a registry layer that enables agents to identify relevant subgraphs, reason about access, and interact with PKGs in a manner that is both efficient and respectful of personal boundaries.

Taken together, insights from autobiographical memory research, life course sociology, and e-government practice converge on a single conclusion: *life events are a natural, cognitively grounded, socially recognised,*

Table 1
Top-level life-event categories and supporting literature

Life-Event Category	Representative Supporting Literature
Identity	Autobiographical memory and identity theory ([4, 12, 13])
Health	Life course theory; institutional life domains; e-government life-event models [6, 9, 10]
Education	Cultural life scripts; life course transitions[3, 7]
Finance	Life course sociology; institutional role transitions [6, 9]
Relationships (incl. family)	Linked lives in life course theory; autobiographical memory [6, 5]
Location (offline and online)	Life events and transitions (e.g. migration); e-government life-event framing [3, 11]
Values & Belief	Identity theory; privacy and sensitive attributes [13, 14]

and institutionally meaningful unit for structuring personal information. In the next section, we draw these threads together and develop a psychologically grounded ontology structure that operationalises life events as the organising substrate for PKG-agent interaction.

2. Psychologically Grounded PKG Ontology Structure

We present a high-level ontology design informed by the psychological model of autobiographical memory and working self of Conway & Pleydell-Pearce's Self-Memory System (SMS) [4], and the sociological model of life course/script, such as Berntsen & Rubin's Cultural Life Script model [5] and the associated Linked Lives concept of sociality, e.g., [8]. Technically, we are informed by Mason's Unified Personal Index and Memory Anchors framework [15]. We emphasise **events**, **contextual anchors**, **identity**, and **goal structures**, which collectively mirror how autobiographical memory is encoded, organised, and retrieved.

This framing establishes the conceptual underpinnings for our model. We now translate these psychological and sociological insights into a concrete *Life Events Registry Ontology* to serve as the semantic backbone for agent-PKG interaction.

2.1. Life Events Registry Ontology

Life events form the central organising backbone of our proposed PKG registry layer. Drawing jointly on the Self-Memory System (SMS) from cognitive psychology [4], cultural life script theory [5], and life course sociology [6, 9], we treat life events not merely as episodic occurrences but as the principal semantic units through which individuals structure personal meaning, institutions classify responsibilities, and agents can effectively scope PKG interaction. In both the SMS and the life course tradition, individuals experience their lives not as undifferentiated flows of data but as sequences of socially recognisable transitions—starting school, moving home, forming relationships, falling ill, changing jobs—which are embedded in culturally shared expectations and institutional systems. These events act as retrieval anchors in autobiographical memory, as turning points in personal identity, and as gateways to institutional processes.

2.1.1. Provenance and Trust in LifeEvent Registries

In agent-mediated PKGs, autonomous reasoning depends not only on the presence of information but on an agent's ability to assess its reliability, recency, and institutional grounding. We therefore treat provenance as a first-class concern for LifeEvent registries.

LifeEvents may be annotated with lightweight provenance sufficient for agent-level trust calibration. At a minimum, this includes: (i) attribution (who asserted or inferred the event), (ii) derivation (self-reported, institutionally certified, or agent-inferred), and (iii) temporal validity (when last verified).

This model aligns naturally with PROV-O, allowing LifeEvents to be associated with generating agents, source documents, or institutional processes. Provenance enables a CUPA to distinguish, for example, between self-asserted educational intentions and formally accredited qualifications, without exposing underlying sensitive records.

Importantly, this provenance layer supports trust-aware agent reasoning while remaining consistent with privacy-by-design principles: fine- grained evidential traces remain local, while abstracted trust signals guide inter-agent coordination.

For agent-mediated PKGs, life events offer three advantages. First, they correspond to cognitively natural divisions of personal experience, enabling the CUPA to reason locally in terms aligned with how people themselves conceptualise their lives. Second, they match the structure of institutional domains: healthcare, education, employment, finance, and social services all operate around life-event triggers rather than raw personal attributes. Third, representing personal data through abstracted life events supports our privacy-by-design strategy: the CUPA can translate detailed, sensitive information into higher-level event-based capability descriptions suitable for secure cooperation with institutional systems. While of course in practice these cooperative interactions will need to be spelled out, for the current purposes, we see policy and legal factors as being largely orthogonal.

To support this, we define a minimal but expressive set of life-event primitives. These are not domain ontologies, but cross-domain scaffolds that provide the semantic hooks from which domain-specific events (e.g. diagnoses, relocations, enrolments) can be modelled. The aim is to capture the shared structural properties for both human cognition and agentic reasoning over life events.

2.1.2. Life Events

To support reasoning and interoperability while preserving privacy, we define *LifeEvent* as a unified abstraction spanning the granularity required by the SMS and life course theory. A LifeEvent may represent an extended transitional period, a cluster of related experiences, or a highly specific occurrence—depending on the user's needs and the CUPA's role. Three subcategories emerge from autobiographical memory theory: *LifetimePeriod* — broad autobiographical intervals (e.g. "my first two years at university", "period of long-term illness") that frame durable commitments, sustained constraints, and ongoing institutional participation; *GeneralEvent* — mid-level groupings of related experiences (e.g. "job searching", "treatment regimen", "relocation process") that capture processes relevant to agent tasks such as planning, recommendations, and eligibility checking; and *EventSpecificDetail (ESK)* — fine-grained perceptual or contextual details that are essential for human recollection but typically remain private to the CUPA, supporting local interpretation without external disclosure.

In addition, life events are embedded in contextual structures—*Memory Anchors*—that situate them temporally, spatially, socially, and environmentally—reflecting the SMS view that memory is cue-dependent and context shapes accessibility. Anchors provide abstracted context (e.g. time-of-year patterns, high-level location, social or institutional roles) to agent reasoning while avoiding exposure of private data. LifeEvents may be annotated with provenance (e.g., `prov:generatedAtTime`, `prov:wasDerivedFrom`, `prov:wasAttributedTo`) to support trust, data quality assessment, and institutional reliability.

Finally, life events are linked to the individual's *PersonalGoals*, which determine the meaning, salience, and implications of events within the working self. Goals give structure to both autobiographical encoding and future-oriented planning; they allow the CUPA to derive high-level capability constraints from detailed event histories. Cultural life scripts provide an optional interpretive scaffold, enabling the CUPA to anticipate sequences of related events or identify domain transitions without imposing prescriptive assumptions.

Together, these components—LifeEvents, their specialisations, contextual anchors, and associated goals—provide a necessary and sufficient basis for the domain-specific life-course categories that follow. They offer a unified, cognitively plausible, institutionally recognisable, and privacy-preserving foundation upon which agent-mediated PKG reasoning can reliably operate.

Having introduced the unified `LifeEvent` abstraction and its grounding in autobiographical memory

and life course theory, we now show how this abstraction manifests across major domains of personal experience. These domains—Identity, Health, Education, Finance, Relationships, Location, and Values & Belief—represent the institutionally and cognitively salient areas in which life events occur. Each subsection illustrates how domain-specific LifeEvents, associated goals, and contextual anchors enable the CUPA to reason privately while exporting only abstract capability signals to other agents.

2.1.3. Aspects of the Life Course

Identity According to [4], a person holds hierarchies of goals, ideal selves, and socially expected selves. The *working self* is a dynamic, goal-driven control system that represents the instantaneous current identity constructed from these: the individual's current projects, values, and self-concept. The working self regulates thinking, behaviour, and memory, sitting between long-term autobiographical knowledge and ongoing cognition, selecting what is relevant, and shaping what can be recalled in relation to current self-beliefs or goals. *Goals* are central to the SMS—they structure encoding and retrieval, and act as "templates" against which new experiences are evaluated. Memories are more accessible when they align with current goals, especially those concerning identity, relationships, competence, and meaning. That is to say, goals drive what gets remembered, how memories are organised, and what is retrievable at a given moment. Access to autobiographical memory is not automatic—it is mediated by the working self and its goals. *Memory* is a goal-modulated construction process, not a passive store.

While of course the psychological self has many elements and behaviours inapplicable to a digital system, these concepts are nonetheless an interesting model for challenges of privacy and disclosure in digital identity. Personal information sharing is highly contextual, and depends on goals, trust, and values, and *interpersonal* information sharing even more so. Having a cognitively natural way to express the different selves and identity facets we share with different audiences may prove to be more effective than a long list of cookie permissions (admittedly, a low bar). These could be particularly powerful in conjunction with, e.g., self-sovereign identities, with the added potential to reduce data linkage.

We propose modelling (life) goals, at any level of generality, in terms of elements of a person's current individual life course script, and as such modelling of personal identity and possible selves can subsist in representations of specific goals and beliefs related to specific life events, relating to health, education, relationships, finance, location, and belief. As such, it makes sense to be non-prescriptive and very minimal in identity-specific vocabularies. Personally-identifying factual information is well covered by ontologies such as FOAF and others. We couple these with a concept of Self, as a container for personal identifying information, events (from the life course or of course prior history), and goals (relating to the life events, and types of event, at some level of granularity). Multiple Self objects can then provide a basis to represent identity facets, self-schemas, turning points events (e.g., major change), and possible selves. Relevant external vocabularies include schema.org (e.g., `schema:Role`, `schema:Occupation`), LODE (e.g., `lode:Event`), and PROV-O (e.g., `prov:generated`).

Health Health is a central domain of the life course and a major organising structure in autobiographical memory. Research in both cognitive and life-course theory shows that health-related events—diagnoses, treatments, accidents, chronic illness trajectories, recoveries—are among the most salient and identity-relevant events in a person's life [6, 9]. Within the SMS [4], health events occupy a special position: they frequently function as *turning points* that reshape the goals, constraints, and capacities of the working self, and consequently reorganise access to autobiographical memory. Illness episodes also tend to be richly encoded with contextual detail, and may deviate from or disrupt culturally expected life scripts, which enhances their accessibility and emotional significance [3, 5].

From a PKG and privacy perspective, health information requires particular care. Health data is among the most sensitive personal data and is strongly protected in institutional contexts (e.g. healthcare, insurance, universities, workplaces). At the same time, agentic systems often need only *functional* insight—e.g. accommodation requirements, accessibility constraints, time-budget constraints during treatment—rather than raw medical detail. This makes the domain ideally suited to our method of "data

minimisation via conceptual abstraction", enabling the CUPA to reason locally over health-related life events and expose only high-level capability statements that respect sensitive boundaries.

We therefore propose modelling health not through a detailed clinical ontology, but via a small vocabulary of health-related life events (e.g. diagnosis events, treatment episodes, recovery periods, accommodation-relevant states) and associated goals (e.g. rehabilitation goals, energy-management goals). Where interoperability is needed, mappings to more detailed vocabularies may be provided. Relevant external ontologies include `schema:MedicalEvent`, `schema:MedicalCondition`, `schema:MedicalProcedu` HL7 FHIR resources (Condition, Observation, Procedure), and SNOMED CT for controlled terminology. By keeping the PKG layer abstract and event-based, and grounding semantics in the life course, we support agent action while maintaining stringent privacy protection.

Education Educational experience is one of the most widely shared and culturally structured domains in the life course. Cultural life script research shows that transitions such as starting school, completing secondary education, entering higher education, or undertaking reskilling later in life reliably appear as normative expectations that shape autobiographical organisation [3]. In the SMS, educational events often function as both identity-relevant turning points and major goal-setting contexts: enrolment in a course or programme typically establishes a hierarchy of learning goals, temporal commitments, and role expectations that influence what the working self treats as salient or feasible.

From a PKG and privacy perspective, educational data is heterogeneous and institutionally sensitive—spanning enrolments, grades, qualifications, course activities, and informal learning. Yet most agent-mediated tasks do not require access to this full record. Instead, they rely on high-level educational LifeEvents (e.g. "currently enrolled in part-time study", "completed prerequisite X", "undertaking professional training") that can be used to infer constraints on availability, recommend appropriate resources, or determine eligibility for institutional services. These abstractions allow the CUPA to privately reason over detailed learning histories while exposing only goal-relevant capability statements to other agents.

We therefore treat education as a domain populated by a family of LifeEvents: enrolment events, study periods, assessment milestones, qualification events—each linked to associated goals (e.g. mastering a skill, completing a credential, preparing for progression) and anchored in temporal, social, and institutional context. This event-based representation accommodates both formal learning (schools, universities, training providers) and informal or self-directed learning, while aligning with life-course expectations and cognitive organisation.

Where explicit interoperability with educational systems is required, domain events may be mapped to existing vocabularies such as the European Learning Model (ELM), which provides structured descriptions of learning activities, qualifications, and credentials. However, within the PKG registry layer, we emphasise minimal abstraction rather than complete modelling: educational LifeEvents serve primarily as navigational and interpretative units for agent reasoning rather than as detailed credential records.

Finance Financial life events are another core institutional domain within life-course research, structuring transitions such as entering paid work, undertaking major expenditures, taking on debt, experiencing income shocks, or planning for retirement [6, 9]. These events reflect socially recognised milestones with clear institutional boundaries—employment systems, banking, taxation, welfare, housing—and therefore map naturally onto an agent-friendly PKG registry. In autobiographical memory, financial transitions often appear as components of broader general events (e.g. starting a first job, moving home, starting a business), anchoring goal-directed reasoning around stability, opportunity, and constraint.

In the Self-Memory System, financial circumstances influence the *working self's* feasible goal set: for example, whether a learner can take on additional study commitments, requires part-time work, or needs budgeting support. However, these considerations do not require disclosing granular financial data to institutional systems. Instead, "capability conditions"—such as availability windows, financial flexibility,

or eligibility for certain opportunities—can be expressed by the CUPA as high-level abstractions of underlying financial life events.

Accordingly, we propose modelling finance at the life-event level: employment starts and ends, income-relevant role transitions, major commitments, and periods of stress or stability. These can be complemented by goals (e.g. saving for a qualification, repayment plans, career aspirations) that allow the CUPA to infer constraints or opportunities without exposing sensitive numbers. External vocabularies that may be mapped where necessary include `schema:JobPosting`, `schema:EmployeeRole`, `schema:MonetaryAmount`, and relevant financial ontologies from the W3C Financial Industry Business Ontology (FIBO). The abstraction-first modelling approach ensures agent usefulness and safeguards financial privacy.

Relationships Relationships—including family ties, friendships, caregiving roles, cohabitation, and partnership histories—constitute one of the strongest organising forces in the life course. Life course theory emphasises the principle of *linked lives*, whereby individual trajectories are interdependent with those of significant others [6, 8]. Autobiographical memory theories likewise show that relational events (forming a partnership, becoming a parent, separation or bereavement) serve as highly accessible memory anchors and often correspond to culturally scripted expectations [5]. In the SMS, relational commitments strongly influence the goals and self-schemas available to the working self.

For PKG-agent interaction, relationship events provide crucial context for workload management, availability estimation, relocation, caregiving responsibilities, and social commitments. Yet these are also deeply sensitive and may involve third-party data. As with identity and health, the goal is therefore not to represent relationship information exhaustively, but to support abstract life-event categories that the CUPA can use privately (e.g. "has childcare responsibilities", "cohabits", "has dependent caregiving role", "recent bereavement"), exporting only agent-relevant implications rather than details.

We propose representing relationships through a minimal vocabulary of relational life events (union, separation, birth/adoption, bereavement, caregiving onset) and relational roles. Relevant terms include `relationship:Relationship`, `foaf:Person`; `schema:parent`, `schema:children`, `schema:spouse` for basic family roles, along with `lode:Event` for relational turning points. Attachment or relational preferences may be represented as `skos:Concepts`.

Location Location—both physical and digital—constitutes a fundamental structuring principle in the life course and in autobiographical memory. Life course research highlights that spatial transitions such as migration, relocation for study or work, household moves, or shifts between online communities represent major life events governed by institutional processes (e.g. housing systems, immigration regimes, campus residency policies) and socially recognised turning points [6, 11]. In the cultural life script literature, moves and relocations routinely appear as salient general events that anchor the organisation of remembered experience [3]. Within the Self-Memory System, location serves not only as a contextual cue for episodic recall, but also as a constraint or enabler of the working self's active goals—for example, determining feasible commute times, study scheduling, social opportunities, or access to services.

In agent-mediated PKGs, location information is both powerful and highly sensitive. Precise location traces, routing histories, or geospatial metadata constitute extremely privacy-invasive data that should not be shared with institutional systems. However, many agent tasks—timetabling, feasibility checking, personalised recommendation, or resource routing—require only *abstracted* knowledge of location-related capabilities (e.g. "available on campus three days per week", "lives in UTC±1 timezone", "relocating soon", "has limited travel range", "participates primarily in online communities"). These correspond closely to location-relevant life events and transitions, allowing the CUPA to reason privately over detailed geospatial context while exporting only goal-relevant implications.

Accordingly, we propose modelling location via a small set of life-event abstractions: residence changes, commutes or travel periods, study abroad or remote work phases, migration events, and digital-community relocations. These can be associated with appropriate *Memory Anchors*—temporal,

spatial, and social—without exposing raw coordinates or continuous tracking data. When optional interoperability is needed, mappings can be provided to vocabularies such as `schema:Place`, `geo:SpatialThing`, and `locn:Geometry`, but the PKG registry layer remains intentionally minimal: it captures stable, agent-interpretable location commitments without surrendering fine-grained location traces. This event-based, privacy-preserving framing aligns with both cognitively natural life segmentation and the operational needs of agent coordination.

Values and belief Values, belief systems, and worldviews play a foundational role in shaping the working self's goal hierarchy and the individual's interpretation of life events. Identity theory highlights the importance of moral commitments, ideological orientations, religious traditions, and meaning-making frameworks in sustaining personal coherence across the life course [13, 12]. Autobiographical memory research similarly notes that value-laden events—conversion, disaffiliation, ritual participation, crises of meaning—often function as deep narrative anchors and identity turning points.

In the SMS, beliefs and values directly modulate what the working self treats as salient, appropriate, or ethically permissible, thereby influencing both memory retrieval and agent-relevant preferences. For example, learners may have restrictions on resource use, working patterns, dietary needs, or scheduling based on religious practice or personal ethics. As with other sensitive domains, institutional systems do not require access to the underlying beliefs—only to actionable implications communicated by the CUPA (e.g. "not available during these hours", "requires specific accommodation", "prefers not to engage with certain content types").

We therefore model values and belief via a small set of life events (ritual participation, value-commitment events, crises of meaning, affiliation or disaffiliation) together with one or more *Self* facets representing value-related identity commitments. External vocabularies that support optional mapping include `schema:Religion`, `skos:ConceptScheme` for belief systems and value taxonomies, `lode:Event` for rituals or conversion events, and `foaf:Group` for communities. As with other domains, the PKG registry layer remains abstract, agent-friendly, and privacy-preserving: the CUPA acts as the interpreter of values, exposing only what is required for ethical agent behaviour.

2.2. Example

To illustrate the LifeEvent registry ontology may be instantiated in practice, Listing 2.2 presents an indicative, non-exhaustive pattern covering the `LifeEvent` primitive, its key specialisations, and an example of a domain-specific event. The listing is intended to be suggestive rather than prescriptive, highlighting how the registry layer provides a lightweight, interoperable structure for CUPA-internal reasoning and controlled exposure to other agents.

Listing 1: Illustrative LifeEvent-based PKG Registry Pattern

```
@prefix pkg: http://pkg.org/model/ .
@prefix schema: http://schema.org/ .
@prefix time: http://www.w3.org/2006/time# .
@prefix prov: http://www.w3.org/ns/prov# .
@prefix lode: http://linkedevents.org/ontology/ .

## rdfs:Class type statements omitted for brevity

## Generic LifeEvent primitive
pkg:LifeEvent rdfs:label "LifeEvent" ;
    rdfs:comment "Abstract event forming the basis of a life-course." .

## Specialisations from autobiographical memory theory
pkg:LifetimePeriod rdfs:subClassOf pkg:LifeEvent ;
    rdfs:comment "Extended period with coherent theme or context." .
```

```
pkg:GeneralEvent rdfs:subClassOf pkg:LifeEvent ;
    rdfs:comment "Cluster_of_occurrences_related_to_goals_or_processes." .

pkg:EventSpecificDetail rdfs:subClassOf pkg:LifeEvent ;
    rdfs:comment "Fine-grained_experiential_detail." .

pkg:StudyPeriod rdfs:subClassOf pkg:LifetimePeriod ;
    rdfs:comment "A_period_of_structured_learning." .

pkg:myMScPeriod a pkg:StudyPeriod ;
    time:hasBeginning "2024-09-01" ;
    time:hasEnd "2026-06-01" ;
    prov:wasAssociatedWith pkg:myLearningGoal .

pkg:myMScPeriod a pkg:StudyPeriod ;
     time:hasBeginning "2024-09-01" ;
     time:hasEnd "2026-06-01" ;
     prov:wasAssociatedWith pkg:myLearningGoal ;
     prov:wasDerivedFrom pkg:OUStudentRecord ;
     prov:generatedAtTime "2024-09-01"^^xsd:date ;
     prov:wasAttributedTo pkg:CUPA .

pkg:myLearningGoal a schema:Action ;
    schema:description "Complete_MSc_qualification" .
```

This mapping demonstrates that a psychologically aligned autobiographical ontology integrates cleanly with widely adopted semantic web vocabularies. The result is an ontology that is both **cognitively realistic** and **interoperable** within Linked Data ecosystems.

3. Agent Interaction and Access Control

3.1. Competency Questions

The following competency questions guided the design of the Life Events Registry Ontology and provide an initial basis for validation:

CQ1: Which current life events constrain availability for work activities?
CQ2: Which institutional agents may interact with the PKG given the current active identity facets?
CQ3: Which life events justify requests for accommodation or support?
CQ4: Which personal goals are active within a given life period?
CQ5: Which capability signals may be safely shared for eligibility checking?
CQ6: Which events are institutionally certified versus self-asserted?
CQ7: Which contextual anchors are relevant for a given agent query?

These questions emphasise relational and goal-oriented properties over deep class hierarchies, supporting agent reasoning while minimising over-engineering.

We propose that representing PKG data in terms of LifeEvents and their associated goals and anchors would enable agents to perform rich, private inference locally, in a manner which fits more closely to human cognitive models. By exporting only the minimal capability signals needed for an external agent to act, we can align with the principle of *data minimisation via conceptual abstraction*: detailed personal data remains local, while high-level, event-derived constraints that guide compliant and context-appropriate actions can be shared in accordance with the relevant working self for any particular context.

From a legal and regulatory perspective, the proposed approach aligns closely with contemporary data protection frameworks such as the GDPR. By abstracting detailed personal data into high-level LifeEvents and capability signals, the model operationalises data minimisation, purpose limitation, and user agency by default. The CUPA acts as an interpretive intermediary, ensuring that only contextually

justified, goal-aligned information is disclosed. As such, LifeEvent registries provide a structural mechanism for lawful, ethical, and explainable agent-mediated data sharing rather than bypassing legal considerations.

4. Conclusion

From our work on developing PKG-based personal stores supporting AI agents assisting students, we know that the first problem one encounters is 'How can an agent explore and use a PKG in a meaningful, efficient way whilst preserving privacy?'. This problem actually relates to the original scenario outlined in the initial paper that started the Semantic Web research area [16]. Our goal is to have a *shared epistemology* instantiated as a *toplevel lightweight ontology*.

Given that we are creating a shared epistemology for all users we elected to use *life events* as our starting point. We know that shared narratives play a key role in human behaviour, and indeed human dominance [17], and therefore have based our indicative Agent Registry Ontology on established social and psychology perspectives on how life events are generally conceptualised.

This paper offers conceptual validation of LifeEvents as an organising substrate for agent-friendly PKG registries. Concrete next steps include: (i) refinement and empirical validation of the ontology using the competency questions outlined above; (ii) implementation of a CUPA prototype supporting provenance-aware LifeEvent reasoning; (iii) user studies evaluating interpretability, trust, and perceived privacy benefits; and (iv) deeper investigation of psychologically grounded models of provenance and trust in autobiographical knowledge.

Implementation will require more detailed ontology engineering and validation with users. The annotation of LifeEvents with *data* provenance with, e.g., PROV-O, is a base level of provenance to support trust and reliability for agentic reasoning; we hypothesise, however, that provenance of memory, and trust in its content, in general has potential to be a fascinating and fruitful area of future investigation if we want it to be psychologically-grounded, as it seems unlikely that this would align directly with technical models of provenance and trust. This is therefore a high priority for future work.

In this paper we have argued that life events—grounded in cognitive psychology, life course sociology, and established e-government practice—offer a compelling foundation for an agent-friendly, privacy-preserving registry layer for PKGs. By unifying episodic, extended, and goal-oriented aspects of personal experience under the `LifeEvent` abstraction, the CUPA can reason locally over rich, sensitive information while exposing only the minimal signals needed for external agents to provide personalised and trustworthy services. This work opens a pathway toward PKG ecosystems that are socially and cognitively relevant and realistic, societally and institutionally aligned, and operationally safe.

Acknowledgments

This work was in part influenced by a conversation with Ian Pickup, Pro Vice-Chancellor Students at the Open University and Sir Tim Berners-Lee on personal agents and conversations with colleagues at Dagstuhl Seminar 25051 on Trust and Accountability in Knowledge Graph-Based AI for Self Determination (https://www.dagstuhl.de/seminars/seminar-calendar/seminar-details/25051).

There are no competing interests related to the development of this research or this paper.

Declaration on Generative AI

During the preparation of this work, the authors used MS Copilot for: Drafting content and Abstract drafting. After using these tool(s)/service(s), the author(s) reviewed and edited the content as needed and take(s) full responsibility for the publication's content.

References

[1] P. A. Bonatti, J. Domingue, A. L. Gentile, A. Harth, O. Hartig, A. Hogan, K. Hose, E. Jimenez-Ruiz, D. L. McGuinness, C. Sun, et al., Towards computer-using personal agents, arXiv preprint arXiv:2503.15515 (2025).

[2] E. Mansour, A. V. Sambra, S. Hawke, M. Zereba, S. Capadisli, A. Ghanem, A. Aboulnaga, T. Berners-Lee, A demonstration of the solid platform for social web applications, in: Proceedings of the 25th international conference companion on world wide web, 2016, pp. 223–226.

[3] D. Berntsen, D. C. Rubin, Cultural life scripts structure recall from autobiographical memory, Memory & Cognition 32 (2004) 427–442.

[4] M. Conway, C. Pleydell-Pearce, The construction of autobiographical memories in the self-memory system, Psychological review 107 (2000) 261–88. doi:10.1037//0033-295X.107.2.261.

[5] D. C. Rubin, D. Berntsen, M. Hutson, The normative and the personal life: Individual differences in life scripts and life story events among usa and danish undergraduates, Memory 17 (2009) 54–68. doi:10.1080/09658210802541442, pMID: 19105087.

[6] G. H. Elder, The life course as developmental theory, Child Development 69 (1998) 1–12.

[7] R. A. Settersten, Age structuring and the rhythm of the life course, in: J. T. Mortimer, M. J. Shanahan (Eds.), Handbook of the Life Course, Springer US, Boston, MA, 2003, pp. 81–98. doi:10.1007/978-0-306-48247-2_4.

[8] R. A. Settersten Jr., Relationships in time and the life course: The significance of linked lives, Research in Human Development 12 (2015) 217–223. doi:10.1080/15427609.2015.1071944.

[9] K. U. Mayer, New directions in life course research, Annual Review of Sociology 35 (2009) 413–433.

[10] F. Sanati, J. Lu, Life-event modelling framework for e-government integration, Electronic Government, an International Journal 7 (2010) 183–202. doi:10.1504/EG.2010.030927.

[11] M. Janssen, G. Kuk, E-government business models for public service networks, International Journal of Electronic Government Research (IJEGR) 3 (2007) 54–71.

[12] A. Giddens, Modernity and self-identity, in: Social theory re-wired, Routledge, 2023, pp. 477–484.

[13] C. Taylor, Sources of the Self, Cambridge, 1989.

[14] D. J. Solove, Understanding privacy, Harvard university press, 2010.

[15] W. Mason, Indaleko: The Unified Personal Index, Ph.D. thesis, The University of British Columbia, 2026. doi:10.48550/arXiv.2602.20507.

[16] T. Berners-Lee, J. Hendler, O. Lassila, Web semantic, Scientific American 284 (2001) 34–43.

[17] Y. N. Harari, Sapiens: A brief history of humankind, Random House, 2014.