

Complements Witness Consistency (Short Paper)

James McKinna
LFCS, School of Informatics
University of Edinburgh
firstname.lastname@ed.ac.uk

Abstract

Much of the existing bx literature, especially that from the PL community on lenses, has described extensional, state-based formalisms. More recently, attention has turned to incorporating intensional information about edits (typically based on monoid actions), or more generally, deltas (typically based on categories), describing how models are updated. Pervasive in both the conceptual modelling, and the mathematics, of varieties of such bx, is the role played by the complement, which generalises the ‘constant complement’ case of the view-update problem in databases. Complements typically reify, or correspond to, data which is abstracted away by passing from a source to a view. In this paper, we present an alternative perspective, which has perhaps been implicit in the lens literature, but not, to our knowledge, previously made explicit anywhere: namely that *elements of the complement are witnesses to the consistency relation maintained by the transformation*. We illustrate this idea with examples drawn from the bx literature, especially that on lenses.

1 Introduction

This short paper can be seen as an attempt to formalise an extended suite of *observations*, which are implicit in the bx literature (though they may be well-known to cognoscenti), especially that on lenses: namely that the (auxiliary) datastructures known as *complements*, witness instances of an underlying consistency relation, itself often left implicit. Moreover, that the lens operations, and the laws they are intended to satisfy, are precisely in the service of witnessing that the lens operations involved do indeed restore consistency after a model change.

Our starting point is the by-now familiar framework, due originally to Meertens [Mee98] and formalised by Stevens [Ste10] in her analysis of the OMG QVT-R standard, of bx considered as forward and backward transformations to restore a consistency relation $R(a, b)$ between elements a, b of model spaces A, B respectively.

In the search for a unifying framework for bx, in which such consistency relations play a *computational* rôle, the paper introduces a definition, that of *generalised complement* below, which uses the ideas of abstract realisability [Läu70], to address the questions

What are *proofs* of consistency? Where/How may they be treated in a computational theory of bx?

Copyright © by the paper’s author. Copying permitted for private and academic purposes.

In: A. Anjorin, J. Gibbons (eds.): Proceedings of the Fifth International Workshop on Bidirectional Transformations (Bx 2016), Eindhoven, The Netherlands, April 8, 2016, published at <http://ceur-ws.org>

The definition offers the slightly more permissive view (as hinted at already in the author’s 2013 Banff BX workshop talk [McK13]), that elements c of such (generalised) complements are *proxies* for (mathematical) proofs of propositions of the form “models $a : A, b : B$ are consistent”. To distinguish such values c from actual proofs, here they are instead dubbed *witnesses to consistency* (the analogy with existing notions of witness structure [CGMS15, for example] is deliberate). In the PL literature on lenses, at least, as illustrated by the examples below, such witnesses arise precisely as elements of the lens *complement*.

2 Abstract Realisability: proofs as first-class citizens

A fundamental idea in the development of mathematical logic in the 20th century is that a proposition should be identified with *the type of its proofs*, popularised as the “Curry-Howard-de Bruijn” correspondence [NG14, for example]. From the point of view of ordinary mathematics, a proposition P holds if and only if we have a proof of it, and we do not care to distinguish between different proofs. This view accords with that of classical logic, where propositions are either true, or false, and hence the set of *proofs* of P may be identified with a singleton set (if P is true), or the empty set (in case P is false). Indeed, one may take a distinguished singleton set, $\mathbf{1} =_{\text{def}} \{*\}$, as standing once-and-for-all as a representative of such singleton sets; such an interpretation of propositions in terms of sets is nowadays called *proof irrelevant*. Indeed, even in the categorical reconstruction of set theory (along intuitionistic lines) [LS88] carried out by Lawvere, Tierney and others in the 1960s and 1970s, a similar move is made, identifying a proposition with a *subset* of $\mathbf{1}$, namely $P \simeq \{* : \mathbf{1} \mid P\}$.

By contrast, since Kleene’s pioneering interpretation [Kle45] of intuitionistic logic in terms of sets of natural numbers (encoding recursive functions), and Kripke’s poset-of-possible-worlds interpretation of modal logic [Kri63], there has been a separate tradition of *proof relevant* interpretations of logic. Läuchli’s framework of *abstract realisability* [Läu70] ushered in the modern era, defining interpretations by identifying, for each well-formed formula P , a set $\text{Prf}(P)$ of abstract *realisers* r of P , which one may think of as candidate, or potential, proofs of P , together with a relation, r *realises* P , such that the proposition P is then identified as that subset of candidate proofs which do in fact realise, or *witness*, that the proposition holds:

$$P \simeq \{r : \text{Prf}(P) \mid r \text{ realises } P\}$$

3 Lens Complements as sets of abstract realisers

To motivate the eventual definition of generalised complement, let us begin with Barbosa et al.’s alternative reformulation [BCF⁺10] of the by-now familiar definition of (very-well-behaved) *asymmetric lens*. In the interests of brevity in what follows, many technical details and proofs have been omitted.

3.1 Asymmetric lens: the very-well-behaved case

An *asymmetric lens* [FGM⁺07] between *non-empty* source type S and view V , with complement C , may be given by the following data: a ‘get’ function $\text{get} : S \rightarrow V$. a ‘put’ function $\text{put} : V \times C \rightarrow S$, and a ‘residual’ function $\text{res} : S \rightarrow C$, satisfying the following properties: ‘GetPut’, $\text{put}(\text{get } s, \text{res } s) = s$; ‘PutGet’, $\text{get}(\text{put}(v, c)) = v$; and ‘PutRes’, $\text{res}(\text{put}(v, c)) = c$. The consistency relation enforced by such a lens is given by $R(s, v) =_{\text{def}} v = \text{get } s$. The PutGet law then states that consistency is indeed restored after an update, i.e. $R(\text{put}(v, c), v)$.

What, then, of the complement set C ? We argue that the element $c = \text{res } s$ is a (indeed: the *unique*) witness to the consistency of s and $\text{get } s$. Accordingly, consider the sets

$$T(s, v) =_{\text{def}} \{\text{res } s \mid R(s, v)\}, \text{ for each } s : S, v : V$$

Then, any inhabitant of $T(s, v)$ establishes the corresponding proposition that $R(s, v)$ holds. With S and V non-empty, PutRes implies that res is surjective, and hence that the image of res , viz. $\bigcup_{s, v} T(s, v)$, is equal to C . Thus, every $c : C$ arises as a witness $c = \text{res } s$ to the corresponding proposition $R(s, \text{get } s)$ for some $s : S$.

3.2 A definition: generalised complement

Given sources A, B , subject to a consistency relation $R \subseteq A \times B$, say that a set C , together with a family of sets $T(a, b)$ (for $a : A, b : B$) forms a **generalised complement** for the relation R , if and only if:

$$(\dagger) \forall a : A, b : B. T(a, b) \text{ inhabited} \quad \text{iff} \quad R(a, b) \quad \text{and, additionally,} \quad (\ddagger) \bigcup_{a : A, b : B} T(a, b) \subseteq C$$

That is to say, C is a set of abstract realisers for R , with $T(a, b)$ delineating those which actually witness $R(a, b)$. **NB.** When A (or B), is *empty*, let $T(a, b) =_{\text{def}} \emptyset$; then *any* C determines a generalised complement for $R =_{\text{def}} \emptyset$.

3.3 Generalised complements for other kinds of lens

3.3.1 Update lens

Ahman and Uustalu have recently considered the following definition [AU14] of *update lens* between source type S and view V , which combines a state-based account of S and V , together with an *edit-based* account of view update, using an edit monoid $\mathbb{P} =_{\text{def}} (P, \circ, \oplus)$: an update lens consists of a ‘lookup’ function $lkp : S \rightarrow V$, an ‘action’ function $act : V \times P \rightarrow V$, and an ‘update’ function $upd : S \times P \rightarrow S$ such that

- ‘action’ is a (right-) action of \mathbb{P} on V , and ‘update’ a (right-) action of \mathbb{P} on S
- ‘lookup’ determines a *homomorphism* between these two actions

(Further details are omitted)

Update lenses do not explicitly identify a notion of lens complement, nor indeed a notion of consistency relation between source and view types. Nevertheless, the foregoing analysis still applies: we may define a consistency relation $R(s, v) =_{\text{def}} v = lkp\ s$. Given any update p to v yielding new view $v' = act(v, p)$, it suffices, in order to restore consistency, to take $s' =_{\text{def}} upd(s, p)$, since lkp is a homomorphism of monoid actions. Furthermore, in particular $R(s, v)$ holds if and only if $\circ \in T(s, v) =_{\text{def}} \{p \mid act(lkp\ s, p) = v\}$, that is, if a distinguished special witness, the identity element \circ , belongs to $T(s, v)$. More generally, consider the relation between s and v defined by $\exists p : P. p \in T(s, v)$, with the monoid element $p : P$ representing the “extent to which s and v may be considered consistent”. For this relaxed notion of consistency, the carrier P of the monoid, together with the sets $T(s, v)$, defines a generalised complement.

Furthermore, since lkp is a homomorphism, if $p \in T(s, v)$, then for any p' we have $p \oplus p' \in T(s, act(v, p'))$, likewise $p' \in T(upd(s, p), act(v, p'))$. Thus, if $R(s, v)$ holds, i.e. $\circ \in T(s, v)$, then $p' \in T(s, act(v, p'))$, and hence also $\circ \in T(upd(s, p'), act(v, p'))$, by straightforward equational reasoning. Thus the more strict notion R of consistency is also restored by the update operations on source and view.

3.3.2 Symmetric lens

Hofmann et al. define a *symmetric lens* [HPW11] between X and Y with complement C in terms of two functions $put_{XY} : X \times C \rightarrow Y \times C$, $put_{YX} : Y \times C \rightarrow X \times C$ subject to certain equational stability laws governing round-tripping (omitted). Then the set C , together with the sets

$$T(x, y) =_{\text{def}} \{c \mid put_{XY}(x, c) = (y, c) \wedge put_{YX}(y, c) = (x, c)\}$$

then defines a generalised complement for the consistency relation $R(x, y) =_{\text{def}} \exists c : C. c \in T(x, y)$, that is, prescribing R in such a way that the conditions (\dagger) and (\ddagger) hold by definition. Then the round-trip laws ensure that given x, y such that $R(x, y)$ holds (witnessed by inhabitants c of $T(x, y)$), and an updated model x' , the operation $put_{XY}(x', c)$ computes an updated model y' , and a new witness $c' \in T(x', y')$, and similarly for put_{YX} .

3.3.3 Edit lens

Technically perhaps the most intricate of the lens constructions developed by Pierce and his collaborators, an *edit lens* [HPW12] between pointed sets (X, x_0) and (Y, y_0) with pointed complement (C, c_0) consists of:

- monoids $\partial X, \partial Y$ of edits, with *partial* actions \bullet_X of ∂X on X and \bullet_Y of ∂Y on Y ;
- stateful monoid homomorphisms $put_{XY} : \partial X \times C \rightarrow \partial Y \times C$ and $put_{YX} : \partial Y \times C \rightarrow \partial X \times C$; and
- a set $K \subseteq X \times C \times Y$, called the *consistency relation*

subject to the conditions (ignoring those governing initial states)

- if $(x, c, y) \in K$, $x' = p \bullet_X x$ is defined for $p \in \partial X$, and $put_{XY}(p, c) = (q, c')$, then $y' = q \bullet_Y y$ is defined and $(x', c', y') \in K$; and the corresponding condition for edits q in ∂Y .

Then, as above, the complement C , together with the sets $T(x, y) =_{\text{def}} \{c \mid (x, c, y) \in K\}$, defines a generalised complement for the relation $R(x, y) =_{\text{def}} \exists c : C. c \in T(x, y)$. The conditions on K indeed then specify that well-defined updates in X , resp. Y , give rise to well-defined updates in Y , resp. X , such that consistency in the sense of R is restored.

3.3.4 Delta lens

Diskin et al. introduced asymmetric and symmetric variants of a notion of *delta lens* [DXC11a, DXC⁺11b], based on the fundamental notion of model space as *category*, generalising previous accounts of updates-via-edit-monoids. Indeed, one can see the (partial) monoid actions $(X, \partial X)$ of a symmetric edit lens as a flattening out of such structure, with homsets $\text{Hom}_X(x, x')$ given by $\{\delta \in \partial X \mid \delta \bullet_X x = x'\}$ etc.

Given categories \mathbb{A}, \mathbb{B} , a *delta lens* between them is then defined in terms of an assignment of *corrs* (for *correspondences*), abstract links r_{ab} relating an object a of \mathbb{A} to object b of \mathbb{B} . The diagram-theoretic properties demanded of such corrs, and their interaction with the structure of identities and composition in \mathbb{A} , resp. \mathbb{B} , amount to the prescription of such links as abstract witnesses to consistency between models a and b , for a generalised complement defined by $T(a, b) =_{\text{def}} \{r_{ab} \mid r_{ab} \text{ is a corr between } a \text{ and } b\}$ and (an implicit) C simply given by the union of all such sets of corrs.

3.4 Other examples

The bx literature is rich in competing approaches, not directly comparable to the approach of the PL literature on lenses. Nevertheless, we consider it fruitful to try to identify candidate structures within such frameworks which may play the role of (generalised) complements, as a basis for further comparison and discussion of the “proof-relevant” perspective on witnessing consistency sketched here.

Indeed, the original starting point for this investigation was an attempt to understand the intricacies of Barbosa et al.’s definition [BCF⁺10] of *matching lens*. There, the structure of the complement plays a dual rôle, enforcing both “chunk consistency”, and for each chunk, consistency with respect to the underlying basic lens. Space considerations forbid further discussion of this example, postponed instead to a future full paper.

4 Conclusions

This paper formalises an observation, which deserves to be more commonplace, in terms of the definition of generalised complement given above, and illustrated through various examples, namely that:

- the consistency relations maintained by bx are witnessed by elements of appropriate auxiliary datastructures;
- such datastructures already play a well-defined role in the formulation, and construction, of such bx, namely in terms of a well-defined generalisation of the notion of (lens) complement.

Such a modest, and perhaps simple-minded, observation, leads us to the conclusion that the designers and implementors of new bx formalisms, or variations thereof, should place front-and-centre of their proposals:

- the consistency relation intended to be maintained by such bx;
- the datastructures which define witnesses to such consistency; and,
- the (obvious) hygiene check that the manipulation of such complements during forward and back restoration steps does, indeed, compute witnesses to such consistency between models being restored.

One may, with good reason, consider that most, if not all, existing bx definitions in the literature fulfil such goals, along the lines sketched in our catalogue of examples. The modest proposal of the current paper is simply that such considerations be made explicit, and thus hopefully provide further insight into the design and implementation of future proposals for bx as computational artefacts.

I speculate that one reason that existing bx/lens definitions do *not* make such data explicit arises from the essentially opportunistic choice of programming language infrastructure, which, at least until recently, has been insufficiently rich to support the notion of types-as-propositions inherent in the definition of generalised complement. Elsewhere [McK16], I intend to pursue the theme of this paper, with a proposal to bring under one foundational roof, in terms of the language and insights from dependent type theory, most if not all existing bx formalisms.

Acknowledgements

This work is funded by EPSRC (EP/K020218/1). I thank the Bx 2016 PC and referees, and my colleagues at Edinburgh and Oxford for their feedback and support during its development.

References

- [AU14] Danel Ahman and Tarmo Uustalu. Coalgebraic update lenses. *ENTCS*, 308:25–48, 2014.
- [BCF⁺10] Davi M.J. Barbosa, Julien Cretin, Nate Foster, Michael Greenberg, and Benjamin C. Pierce. Matching lenses: Alignment and view update. In *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming, ICFP '10*, pages 193–204, New York, NY, USA, 2010. ACM.
- [CGMS15] James Cheney, Jeremy Gibbons, James McKinna, and Perdita Stevens. Towards a principle of least surprise for bidirectional transformations. In *Proceedings of the 4th International Workshop on Bidirectional Transformations co-located with Software Technologies: Applications and Foundations, STAF 2015, L'Aquila, Italy, July 24, 2015.*, pages 66–80, 2015.
- [DXC11a] Zinovy Diskin, Yingfei Xiong, and Krzysztof Czarnecki. From state- to delta-based bidirectional model transformations: the asymmetric case. *JOT*, 10:6: 1–25, 2011.
- [DXC⁺11b] Zinovy Diskin, Yingfei Xiong, Krzysztof Czarnecki, Hartmut Ehrig, Frank Hermann, and Fernando Orejas. From state- to delta-based bidirectional model transformations: the symmetric case. In *MODELS*, volume 6981 of *LNCS*. Springer, 2011.
- [FGM⁺07] J. Nathan Foster, Michael B. Greenwald, Jonathan T. Moore, Benjamin C. Pierce, and Alan Schmitt. Combinators for bidirectional tree transformations: A linguistic approach to the view-update problem. *ACM TOPLAS*, 29(3):17, 2007.
- [HPW11] Martin Hofmann, Benjamin C. Pierce, and Daniel Wagner. Symmetric lenses. In *POPL*, pages 371–384. ACM, 2011.
- [HPW12] Martin Hofmann, Benjamin C. Pierce, and Daniel Wagner. Edit lenses. In *POPL*, pages 495–508. ACM, 2012.
- [Kle45] S.C. Kleene. On the interpretation of intuitionistic number theory. *Journal of Symbolic Logic*, 10(4):109–124, 1945. doi:10.2307/2269016.
- [Kri63] S. Kripke. Semantical Analysis of Modal Logic I. Normal Modal Propositional Calculi. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 9:67–96, 1963.
- [Läu70] H. Läuchli. An abstract notion of realizability for which intuitionistic predicate calculus is complete. In A. Kino, J. Myhill, and R. E. Vesley, editors, *Intuitionism and Proof Theory*. North Holland, 1970.
- [LS88] J. Lambek and P. J. Scott. *Introduction to Higher-Order Categorical Logic*. Cambridge Studies in Advanced Mathematics. CUP, 1988.
- [McK13] James McKinna. Type theory and algebraic bx. Video of a talk given at the 2013 Banff Bx meeting, 2013. <http://videos.birs.ca/2013/13w5115/201312040919-McKinna.mp4>.
- [McK16] James McKinna. Bidirectional transformations with deltas: a dependently typed approach. Talk proposal accepted for Bx 2016. Manuscript in preparation, 2016.
- [Mee98] Lambert Meertens. Designing constraint maintainers for user interaction. Unpublished manuscript, available from <http://www.kestrel.edu/home/people/meertens/>, June 1998.
- [NG14] Rob Nederpelt and Herman Geuvers. *Type Theory and Formal Proof*. CUP, 2014.
- [Ste10] Perdita Stevens. Bidirectional model transformations in QVT: Semantic issues and open questions. *SoSyM*, 9(1):7–20, 2010.